Spectral Approaches to Nearest Neighbor Search

Amirali Abdullah * University of Utah
Alexandr Andoni Microsoft Research
Ravindran Kannan Microsoft Research
Robert Krauthgamer† Weizmann Institute

arXiv:1408.0751v1 [cs.DS] 4 Aug 2014

Abstract

We study spectral algorithms for the high-dimensional Nearest Neighbor Search problem (NNS). In particular, we consider a semi-random setting where a dataset P in $\mathbb{R}^d$ is chosen arbitrarily from an unknown subspace of low dimension $k \ll d$, and then perturbed by fully d-dimensional Gaussian noise. We design spectral NNS algorithms whose query time depends polynomially on d and $\log n$ (where $n = |P|$) for large ranges of k, d and n. Our algorithms use a repeated computation of the top PCA vector/subspace, and are effective even when the random-noise magnitude is *much larger* than the inter-point distances in P. Our motivation is that in practice, a number of spectral NNS algorithms outperform the random-projection methods that seem otherwise theoretically optimal on worst case datasets. In this paper we aim to provide theoretical justification for this disparity.

*A large of this work was carried out and funded at Microsoft Research. This effort was also supported in part by the NSF under grant #CCF-0953066 Email: amirali@cs.utah.edu

†Work done in part while at Microsoft Silicon Valley. Work supported in part by the Israel Science Foundation grant #897/13, the US-Israel BSF grant #2010418, and by the Citi Foundation. Email: robert.krauthgamer@weizmann.ac.il

1 Introduction

A fundamental tool in high-dimensional computational geometry is the random projection method. Most notably, the Johnson-Lindenstrass Lemma [JL84] says that projecting onto a uniformly random k-dimensional subspace of $\mathbb{R}^d$ approximately preserves the distance between any (fixed) points $x, y \in \mathbb{R}^d$ (up to scaling), except with probability exponentially small in k. This turns out to be a very powerful approach as it effectively reduces the dimension from d to a usually much smaller k via a computationally cheap projection, and as such has had a tremendous impact on algorithmic questions in high-dimensional geometry.

A classical application of random projections is to the high-dimensional Nearest Neighbor Search (NNS) problem. Here we are given a dataset of n points from $\mathbb{R}^d$, which we preprocess to subsequently find, given a query point $q \in \mathbb{R}^d$, its closest point from the dataset. It is now well-known that exact NNS admits algorithms whose running times have good dependence on n, but exponential dependence on the dimension d [Mei93, Cla88]; however these are unsatisfactory for moderately large d.

To deal with this "curse of dimensionality", researchers have studied algorithms for *approximate* NNS, and indeed in the high-dimensional regime, many, if not all, of these algorithms rely heavily on the random projection method. Consider the case of Locality-Sensitive Hashing (LSH), introduced in [IM98], which has been a theoretically and practically successful approach to NNS. All known variants of LSH for the Euclidean space, including [IM98, DIIM04, AI06, AINR14], involve random projections.[1] For example, the space partitioning algorithm of [DIIM04] can be viewed as follows. Project the dataset onto a random k-dimensional subspace, and impose a randomly-shifted uniform grid. Then, to locate the near(est) neighbor of a point q, look up the points in the grid cell where q falls into. Usually, this space partitioning is repeated a few times to amplify the probability of success (see also [Pan06]).

While random projections work well and have provable guarantees, it is natural to ask whether one can improve the performance by replacing "random" projections with "best" projections. Can one optimize the projection to use — and the space partitioning more generally — as a *function of the dataset at hand*? For example, in some tasks requiring dimension reduction, practitioners often rely on Principal Component Analysis (PCA) and its variants. Indeed, in practice, this consideration led to numerous successful heuristics such as PCA tree [Spr91, McN01, VKD09] and its variants (called randomized kd-tree) [SAH08, ML09], spectral hashing [WTF08], semantic hashing [SH09], and WTA hashing [YSRL11], to name just a few. Oftentimes, these heuristics outperform algorithms based on vanilla random projections. All of them adapt to the dataset, including many that perform some *spectral* decomposition of the dataset. However, in contrast to the random projection method, none of these methods have rigorous correctness or performance guarantees.

Bridging the gap between random projections and data-aware projections has been recognized as a big open question in Massive Data Analysis, see e.g. a recent National Research Council report [NRC13, Section 5]. The challenge here is that random projections are themselves (theoretically) optimal not only for dimensionality reduction [Alo03, JW13], but also for some of its algorithmic applications [IW03, Woo04], including NNS in certain settings [AIP06]. We are aware of only one work addressing this question: data-dependent LSH, which was introduced recently [AINR14], provably improves the query time polynomially. However, their *space partitions* are very different from the aforementioned practical heuristics (e.g., they are not spectral-based), and do not explain why data-aware *projections* help at all.

In this work, we address this gap by studying data-aware projections for the nearest neighbor search problem. As argued above, for worst-case inputs we are unlikely to beat the performance of random projections, and thus it seems justified to revert to the framework of smoothed analysis [ST09] to study the gap in practical performance. We consider a semi-random model, where the dataset is formed by first taking an arbitrary (worst-case) set of n points in a k-dimensional subspace of $\mathbb{R}^d$, and then perturbing each point by adding to

[1] While [IM98] is designed for the Hamming space, their algorithm is extended to the Euclidean space by an embedding of ℓ_2 into ℓ_1, which itself uses random projections [JS82].

it Gaussian noise $N_d(0, \sigma^2 I_d)$. The query point is selected using a similar process. Our algorithms are able to find the query's nearest neighbor as long as there is a small gap (1 vs $1+\epsilon$) in the distance to the nearest neighbor versus other points in the *unperturbed* space — this is a much weaker assumption than assuming the same for the perturbed points.

Most importantly, our results hold even when the noise magnitude is *much larger* than the distance to the nearest neighbor. The noise vector has length (about) $\sigma\sqrt{d}$, and so for $\sigma \gg 1/\sqrt{d}$, the noise magnitude exceeds the original distances. In such a case, a random Johnson-Lindenstrauss projection to a smaller dimension will not work — the error due to the projection will lose all the information on the nearest neighbor.

In contrast, our results show that data-aware projections provably guarantee good performance in this model. We describe the precise model in Section 2.

1.1 Algorithmic Results

We propose two spectral algorithms for nearest neighbor search, which achieve essentially the same performance as NNS algorithms in k and $O(k \log k)$-dimensional space, respectively. These spectral algorithms rely on computing a PCA subspace or vector respectively — the span of the singular vectors corresponding to the top singular values of an $n \times d$ matrix representing some n points in $\mathbb{R}^d$. Our algorithms are inspired by PCA-based methods that are commonly used in practice for high-dimensional NNS, and we believe that our rigorous analysis may help explain (or direct) those empirical successes. We defer the precise statements to the respective technical sections (specifically, Theorems 5.1 and 6.1), focusing here on a *qualitative* description.

The first algorithm performs *iterative PCA*. Namely, it employs PCA to extract a subspace of dimension (at most) k, identifies the points captured well by this subspace, and then repeats iteratively on the remaining points. The algorithm performs at most $O(\sqrt{d \log n})$ PCAs in total, and effectively reduces the original NNS problem to $O(\sqrt{d \log n})$ instances of NNS in k dimensions. Each of these NNS instances can be solved by any standard low-dimensional $(1+\epsilon)$-approximate NNS, such as [Cla94, AMN$^+$98, AMM09, HP01, AM02, CG06], which can give, say, query time $(1/\epsilon)^{O(k)} \log^2 n$. See Section 5, and the crucial technical tool it uses in Section 4. As a warmup, we initially introduce a simplified version of the algorithm for a (much) simpler model in Section 3.

The second algorithm is a variant of the aforementioned PCA tree, and constructs a tree that represents a recursive space partition. Each tree node corresponds to finding the top PCA direction, and partitioning the dataset into slabs perpendicular to this direction. We recurse on each slab until the tree reaches depth $2k$. The query algorithm searches the tree by following a small number of children (slabs) at each node. This algorithm also requires an additional preprocessing step that ensures that the dataset is not overly "clumped". The overall query time is $(k/\epsilon)^{O(k)} \cdot d^2$. See Section 6.

While the first algorithm is randomized, the second algorithm is deterministic and its failure probability comes only from the semi-random model (randomness in the instance).

1.2 Related Work

There has been work on understanding how various tree data structures adapt to a *low-dimensional* pointset, including [DF08, VKD09]. For example, [VKD09] show that PCA trees adapt to a form of "local covariance dimension", a spectrally-inspired notion of dimension, in the sense that a PCA tree halves the "diameter" of the pointset after a number of levels dependent on this dimension notion (as opposed to the ambient dimension d). Our work differs in a few respects. First, our datasets do not have a small local covariance dimension. Second, our algorithms have guarantees of performance and correctness for NNS for a worst-case query point (e.g., the true nearest neighbor can be any dataset point). In contrast, [VKD09] prove a

guarantee on diameter progress, which does not necessarily imply performance guarantees for NNS, and, in fact, may only hold for average query point (e.g., when the true nearest neighbor is random). Indeed, for algorithms considered in [VKD09], it is easy to exhibit cases where NNS fails.[2]

For our model, it is tempting to design NNS algorithms that find the original k-dimensional subspace and thus "de-noise" the dataset by projecting the data onto it. This approach would require us to solve the ℓ_∞-regression problem with high precision.[3] Unfortunately, this problem is NP-hard in general [GK93], and the known algorithms are quite expensive, unless k is constant. Har-Peled and Varadarajan [HV04] present an algorithm for ℓ_∞-regression achieving $(1+\epsilon)$-approximation in time $O(nde^{e^{O(k^2)}\epsilon^{-2k-3}})$, which may be prohibitive when $k \geq \Omega(\sqrt{\log\log n})$. In fact, there is a constant $\delta > 0$ such that it is Quasi-NP-hard, i.e., implies $\mathrm{NP} \subseteq \mathrm{DTIME}(2^{(\log n)^{O(1)}})$, to even find a $(\log n)^\delta$ approximation to the best fitting k-subspace when $k \geq d^\epsilon$ for any fixed $\epsilon > 0$ [VVZ02].

We also note that the problem of finding the underlying k-dimensional space is somewhat reminiscent of the learning mixtures of Gaussians problem [Das99]; see also [AK01, VW04, MV10] and references therein. In the latter problem, the input is n samples generated from a mixture of k Gaussian distributions with unknown mean (called centers), and the goal is to identify these centers (the means of the k Gaussians). Our setting can be viewed as having n centers in a k-dimensional subspace of $\mathbb{R}^d$, and the input contains exactly one sample from each of the n centers. Methods known from learning mixtures of Gaussians rely crucially on obtaining multiple samples from the same center (Gaussian), and thus do not seem applicable here.

Finally, the problem of nearest neighbor search in Euclidean settings with "effective low-dimension" has received a lot of attention in the literature, including [KR02, KL04, HM05, BKL06, CG06, IN07] among many others. Particularly related is also the work [HK13], where the authors consider the case when the dataset is high-dimensional, but the query comes from a (predetermined) low-dimensional subspace. These results do not seem readily applicable in our setting because our dataset is really high-dimensional, say in the sense that the doubling dimension is $\Omega(d)$.

1.3 Techniques and Ideas

We now discuss the main technical ideas behind our two algorithms. First, we explain why some natural ideas do *not* work. The very first intuitive line of attack to our problem is to compute a k-dimensional PCA of the pointset, project it into this k-dimensional subspace, and solve the k-dimensional NNS problem there. This approach fails because the noise is too large, and PCA only optimizes *the sum of distances* (i.e., an average quantity, as opposed to the "worst case" quantity). In particular, suppose most of the points lie along some direction $\vec{u}$ and only a few points lie in the remaining dimensions of our original subspace U (which we call sparse directions). Then, the k-PCA of the dataset will return a top singular vector close to $\vec{u}$, but the remaining singular vectors will be mostly determined by random noise. In particular, the points with high component along sparse directions may be very far away from the subspace returned by our PCA, and hence "poorly-captured" by the PCA space. Then, the NNS data structure on the projected points will fail for some query points. If the difference between a query point q and its nearest neighbor p^* is along $\vec{u}$, whereas the difference between q and a poorly-captured point p' is along the sparse directions, then any such p' will cluster around q in the k-PCA space, at a distance much closer than $\|q - p^*\|$. See Figure 1.

[2] For example, if we consider the top PCA direction of a dataset and the median threshold, we can plant a query–near-neighbor pair on the two sides of the partition. Then, this pair, which won't affect top PCA direction much, will be separated in the PCA tree right from the beginning.

[3] As we explain later, related problems, such as ℓ_2-regression would not be sufficient.

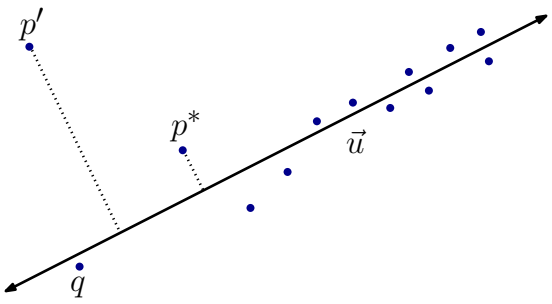

Figure 1: p' lies along a poorly captured direction.

Our first algorithm instead runs k-PCAs *iteratively*, while pruning away points "well-captured" by the PCA (i.e., close to the PCA subspace). In particular, this allows us to discover the points in sparse directions in *later* iterations. Furthermore, to ensure correctness of nearest neighbor queries in presense of large noise, we do not necessarily take all the top k singular values, but only those that exceed some threshold value; this guarantees that all directions in our PCA subspace have a large component inside U. Showing this guarantee analytically is non-trivial, starting even with the definition of what it means to be "close" for two spaces, which may have different dimensions. For this purpose, we employ the so-called $\sin\theta$ machinery, which was developed by Davis and Kahan [DK70] and by Wedin [Wed72], to bound the perturbations of the singular *vectors* of a matrix in presence of noise. Notice the difference from the more usual theory of perturbations of the singular *values*. For example, in contrast to singular values, it is not true that the top singular vector is "stable" when we perturb a given matrix.

The actual algorithm has one more important aspect: in each iteration, the PCA space is computed on a *sample* of the (surviving) data points. This modification allows us to control spurious conditioning induced by earlier iterations. In particular, if instead we compute the PCA of the full data, once we argue that a vector $\tilde{p}$ "behaves nicely" in one iteration, we might effectively condition on the direction of its noise, potentially jeopardizing noise concentration bounds on later iterations. (While we do not know if sampling is really necessary for the algorithm to work, we note that practically it is a very reasonable idea to speed up preprocessing nonetheless.)

The second algorithm is based on the PCA-tree, which partitions the space recursively, according to the top PCA direction. This can be seen as another (extreme) form of "iterative PCA". At each node, the algorithm extracts one top PCA direction, which always contains the "maximum" information about the dataset. Then it partitions the dataset into a few slabs along this direction, thereby partitioning the datasets into smaller parts "as quickly as possible". This allows the tree to narrow down on the sparse directions quicker. The performance of the PCA tree depends exponentially on its depth, hence the crux of the argument is to bound the depth. While it seems plausibly easy to show that a partitioning direction should never be repeated, this would give too loose a bound, as there could be a total of $\approx \exp(k)$ essentially distinct directions in a k-dimensional space. Instead, we perform a mild form of orthonormalization as we progress down the tree, to ensure only $O(k)$ directions are used in total. In the end, the query time is roughly $k^{O(k)}$, i.e., equivalent to a NNS in an $O(k \log k)$-dimensional space.

We note that this algorithm has two interesting aspects. First, one has to use *centered* PCA, i.e., PCA on the data centered at zero: otherwise, every small error in PCA direction may move points a lot for subsequent iterations, misleading a non-centered PCA. Second, from time to time, we need to do "de-clumping" of the data, which essentially means that the data is sparsified if the points are too close to each other. This operation also appears necessary; otherwise, a cluster of points that are close in the original space, might mislead the PCA due to their noise components. Furthermore, in contrast to the first algorithm, we cannot afford to iterate through $\approx d$ iterations to eliminate "bad" directions one by one.

2 The Model

We assume throughout the dataset is generated as follows.[4] Let U be a k-dimensional subspace of $\mathbb{R}^d$. Let $P = \{p_1, \ldots, p_n\}$ be a set of n points all living in U and having at least unit norm, and let $q \in U$ be a query point. We assume that $d = \Omega(\log n)$. The point set P is perturbed to create $\tilde{P} = \{\tilde{p}_1, \ldots, \tilde{p}_n\}$ by adding to each point independent Gaussian noise, and the query point q is perturbed similarly. Formally,

$$\tilde{p}_i = p_i + t_i \text{ where } t_i \sim N_d(0, \sigma I_d), \qquad \forall p_i \in P, \tag{2.1}$$

$$\tilde{q} = q + t_q \text{ where } t_q \sim N_d(0, \sigma I_d). \tag{2.2}$$

Let us denote the nearest neighbor to q in P by p^* and let $\tilde{p}^*$ be its perturbed version. We shall actually consider the *near-neighbor* problem, by assuming that in the unperturbed space, there is one point $p^* \in P$ within distance 1 from the query, and all other points are at distance at least $1 + \epsilon$ from the query, for some known $0 < \epsilon < 1$. Formally,

$$\exists p^* \in P \text{ such that } \|q - p^*\| \leq 1 \text{ and } \forall p \in P \setminus \{p^*\},\ \|q - p\| \geq 1 + \epsilon. \tag{2.3}$$

We note that even if there is more than one such point p^* so that $\|q - p^*\| \leq 1$, our algorithms will return one of these close p^* correctly. Also our analysis in Section 5 extends trivially to show that for any x such that $x \geq 1$ and $\|q - p^*\| = x$, our first algorithm the iterative PCA actually returns a $(1 + \epsilon)$-approximate nearest neighbor to q. We omit the details of this extended case for ease of exposition.

2.1 Preliminary Observations

For the problem to be interesting, we need that the perturbation does not change the nearest neighbor, i.e., $\tilde{p}^*$ remains the closest point to $\tilde{q}$. We indeed show this is the case as long as $\sigma \ll \epsilon / \sqrt[4]{d \log n}$. Notice that the total noise magnitude is roughly $\sigma\sqrt{d}$, which can be much larger than 1 (the original distance to the nearest neighbor). Hence after the noise is added, the ratio of the distance to the nearest neighbor and to other (nearby) points becomes very close to 1. This is the main difficulty of the problem, as, for example, it is the case where random dimensionality reduction would lose nearest neighbor information. We recommend to keep in mind the following parameter settings: $k = 20$ and $\epsilon = 0.1$ are constants, while $d = \log^3 n$ and $\sigma = \Theta(1/\log n)$ depend asymptotically on $n = |P|$. In this case, for example, our algorithms actually withstand noise of magnitude $\Theta(\sqrt{\log n}) \gg 1$.

Here and in the rest of the paper, we will repeatedly employ concentration bounds for Gaussian noise, expressed as tail inequalities on χ^2 distribution. We state here for reference bounds from [LM98], where χ^2_d is the same distribution as $\|N_d(0, I_d)\|_2^2$. The term *with high probability (w.h.p.)* will mean that probability $1 - n^{-C}$ for sufficiently large $C > 0$.

Theorem 2.1. *([LM98]) Let $X \sim \chi^2_d$. For all $x \geq 0$,*

$$\Pr\left[X \geq d\left(1 + 2\sqrt{\tfrac{x}{d}}\right) + x\right] \leq e^{-x}, \quad \textit{and} \quad \Pr\left[X \leq d\left(1 - 2\sqrt{\tfrac{x}{d}}\right)\right] \leq e^{-x}.$$

Corollary 2.2. *For $n \geq 1$, let $X \sim \chi^2_d$. Then $\Pr[|X - d| \geq d + 4\sqrt{d \log n} + 4 \log n] \leq \frac{2}{n^4}$.*

We now show that after the perturbation of P, q, the nearest neighbor of $\tilde{q}$ will remain $\tilde{p}^*$, w.h.p.

Lemma 2.3. *Consider the above model (2.1)-(2.3) for $n > 1$, $\epsilon \in (0, 1)$, dimensions $k < d = \Omega(\log n)$, and noise standard deviation $\sigma \leq c\epsilon / \sqrt[4]{d \log n}$, where $c > 0$ is a sufficiently small constant. Then w.h.p. the nearest neighbor of $\tilde{q}$ (in $\tilde{P}$) is $\tilde{p}^*$.*

[4] An exception is the warm-up Section 3, where the noise is small adversarial.

Proof. Write $(X_1, \ldots, X_d)^{\mathrm{T}} = \tilde{q} - q \sim N_d(0, \sigma I_d)$, and similarly $(Y_1, \ldots, Y_d)^{\mathrm{T}} = \tilde{p}^* - p \sim N_d(0, \sigma I_d)$. Let $Z_i = X_i - Y_i$, and note that the Z_i's are independent Gaussians, each with mean 0 and variance $2\sigma^2$. Then by direct computation

$$\|\tilde{q} - \tilde{p}^*\|^2 = 1 + \sum_{i=1}^{d}(X_i - Y_i)^2 + \sum_{i=1}^{d}(X_i - Y_i)(q_i - p_i^*) = 1 + \sum_{i=1}^{d} Z_i^2 + \sum_{i=1}^{d} Z_i(q_i - p_i^*). \tag{2.4}$$

For the term $\sum_{i=1}^{d} Z_i^2$, Theorem 2.1 gives us $\Pr\left[\left|\sum_{i=1}^{d} Z_i^2 - 2\sigma^2 d\right| \geq 2\sigma^2\left(x + d \cdot 2\sqrt{\frac{x}{d}}\right)\right] \leq 2e^{-x}$. Setting $x = 4\log n \leq O(d)$, observe that $2\sigma^2(x + 2\sqrt{xd}) \leq O(\sigma^2\sqrt{xd}) \leq O(c^2\epsilon^2)$, and thus we have $\Pr\left[\left|\sum_{i=1}^{d} Z_i^2 - 2d\sigma^2\right| \geq O(c^2\epsilon^2)\right] \leq \frac{2}{n^4}$.

Now the term $\sum_{i=1}^{d} Z_i(q_i - p_i^*)$ is a Gaussian with mean 0 and variance $\sum_{i=1}^{d}(q_i - p_i^*)^2 \operatorname{Var}[Z_i] = 2\sigma^2\|q - p^*\|^2 = 2\sigma^2$, and thus with high probability $|\sum_{i=1}^{d} Z_i(q_i - p_i^*)| \leq O(\sigma\sqrt{\log n})$. Substituting for σ and recalling $d = \Omega(\log n)$, the righthand-side can be bounded by $O(c\epsilon)$. Altogether, with high probability

$$\|\tilde{q} - \tilde{p}^*\|^2 \leq 1 + 2d\sigma^2 \pm O(c^2\epsilon^2) \pm O(c\epsilon). \tag{2.5}$$

Similarly, for every other point $p \neq p^*$, with high probability,

$$\|\tilde{q} - \tilde{p}\|^2 \geq \|q - p\|^2 + 2d\sigma^2 \pm O(c^2\epsilon^2) \pm O(c\epsilon)\|q - p\|, \tag{2.6}$$

and we can furthermore take a union bound over all such points $p \neq p^*$. Now comparing Eqns. (2.5) and (2.6) when $\|q - p\|^2 \geq 1 + \epsilon$ and $c > 0$ is sufficiently small, gives us the desired conclusion. □

Remark 2.1. *The problem remains essentially the same if we assume the noise has no component in the space U. Indeed, we can absorb the noise inside U into the "original" points (P and q). With high probability, this changes the distance from q to every point in P by at most $O(\sigma\sqrt{k\log n}) \ll \epsilon$. Hence, in the rest of the article, we will assume the noise is perpendicular to U.*

3 Warmup: Iterative PCA under Small Adversarial Noise

To illustrate the basic ideas in our "iterative PCA" approach, we first study it in an alternative, simpler model that differs from Section 2 in that the noise is *adversarial* but of *small magnitude*. The complete "iterative PCA" algorithm for the model from Section 2 will appear in Section 5.

In the bounded noise model, for fixed $\epsilon \in (0, 1)$, we start with an n-point dataset P and a point q, both lying in a k-dimensional space $U \subset \mathbb{R}^d$, such that

$$\exists p^* \in P \text{ such that } \|q - p^*\| \leq 1 \text{ and } \forall p \in P \setminus \{p^*\},\ \|q - p\| \geq 1 + \epsilon \tag{3.1}$$

The set $\tilde{P}$ consists of points $\tilde{p}_i = p_i + t_i$ for all $p_i \in P$, where the noise t_i is arbitrary, but satisfies $\|t_i\| \leq \epsilon/16$ for all i. Similarly, $\tilde{q} = q + t_q$ with $\|t_q\| \leq \epsilon/16$.

Theorem 3.1. *Suppose there is a $(1+\epsilon/4)$-approximate NNS data structure for n points in a k-dimensional Euclidean space with query time F_{query}, space F_{space}, and preprocessing time F_{prep}. Then for the above adversarial-noise model, there is a data structure that preprocesses $\tilde{P}$, and on query $\tilde{q}$ returns $\tilde{p}^*$. This data structure has query time $O((dk+F_{\mathrm{query}})\log n)$, space $O(F_{\mathrm{space}})$, and preprocessing time $O(n+d^3+F_{\mathrm{prep}})$.*

First we show that the nearest neighbor "remains" p^* even after the perturbations (similarly to Lemma 2.3. Let $\alpha = \epsilon/16$.

Claim 3.2. *The nearest neighbor of $\tilde{q}$ in $\tilde{P}$ is $\tilde{p}^*$.*

Proof. For all i, we have $\|\tilde{p}_i - p_i\| \leq \|t_i\| \leq \alpha$, hence by the triangle inequality, $\|\tilde{q} - \tilde{p}^*\| \leq \|\tilde{q} - q\| + \|q - p^*\| + \|p^* - \tilde{p}^*\| \leq \|q - p^*\| + 2\alpha$. For all $p \neq p^*$, a similar argument gives $\|\tilde{q} - \tilde{p}\| \geq \|q - p^*\| + \epsilon - 2\alpha$. □

We now describe the algorithm used to prove Theorem 3.1. Our algorithm first finds a small collection $\mathcal{U}$ of k-dimensional subspaces, such that every point of $\tilde{P}$ is "captured well" by at least one subspace in $\mathcal{U}$. We find this collection $\mathcal{U}$ by iteratively applying PCA, as follows (see Algorithm 3.1). First compute the top (principal) k-dimensional subspace of $\tilde{P}$. It "captures" all points $\tilde{p} \in \tilde{P}$ within distance $\sqrt{2}\alpha$ from the subspace. Then we repeat on the remaining non-captured points, if any are left. In what follows, let $p_{\tilde{U}}$ denote the projection of a point p onto $\tilde{U}$, and define the distance between a point x and a set (possibly a subspace) S as $d(x, S) = \inf_{y \in S} \|x - y\|$.

Algorithm 3.1 Iteratively locate subspaces

$j \leftarrow 0$; $\tilde{P}_0 \leftarrow \tilde{P}$
while $\tilde{P}_j \neq \emptyset$ **do**
  $\tilde{U}_j \leftarrow$ the k-dimensional PCA subspace of $\tilde{P}_j$
  $M_j \leftarrow \{\tilde{p} \in \tilde{P}_j : \ d(\tilde{p}, \tilde{U}_j) \leq \sqrt{2}\alpha\}$
  $\tilde{P}_{j+1} \leftarrow \tilde{P}_j \setminus M_j$
  $j \leftarrow j + 1$
end while
return $\tilde{\mathcal{U}} = \{\tilde{U}_0, \ldots, \tilde{U}_{j-1}\}$ and the associated point sets $\{M_0, M_1, \ldots, M_{j-1}\}$.

The remainder of the preprocessing algorithm just constructs for each subspace $\tilde{U} \in \mathcal{U}$ a data structure for k-dimensional NNS, whose dataset is the points captured by $\tilde{U}$ projected onto this subspace $\tilde{U}$ (treating $\tilde{U}$ as a copy of $\mathbb{R}^k$). Overall, the preprocessing phase comprises of $O(\log n)$ PCA computations and constructing $O(\log n)$ data structures for a k-dimensional NNS.

The query procedure works as follows. Given a query point $\tilde{q}$, project $\tilde{q}$ onto each $\tilde{U} \in \mathcal{U}$ to obtain $\tilde{q}_{\tilde{U}}$, and find in the data structure corresponding to this $\tilde{U}$ a $(1 + \epsilon/4)$-approximate nearest neighbor point $\tilde{p}_{\tilde{U}}$ to $\tilde{q}_{\tilde{U}}$. Then compute the distance between $\tilde{q}$ and each $\tilde{p}$ (original points corresponding to $\tilde{p}_{\tilde{U}}$), and report the the closest one to $\tilde{q}$.

We now proceed to analyze the algorithm.

Claim 3.3. *Algorithm 3.1 terminates within $O(\log n)$ iterations.*

Proof. Let U be the PCA subspace of P and let $\tilde{U}$ be the PCA subspace of $\tilde{P}$. Since $\tilde{U}$ minimizes (among all k-dimensional subspaces) the sum of squared distances from all $\tilde{p} \in \tilde{P}$ to $\tilde{U}$,

$$\sum_{\tilde{p} \in \tilde{P}} d(\tilde{p}, \tilde{U})^2 \leq \sum_{\tilde{p} \in \tilde{P}} d(\tilde{p}, U)^2 \leq \sum_{\tilde{p} \in \tilde{P}} \|\tilde{p} - p\|^2 \leq \alpha^2 n.$$

Hence, at most half of the points in $\tilde{P}$ may have distance to $\tilde{U}$ which is greater than $\sqrt{2}\alpha$. The current set M will capture the other (at least a half fraction) points, and the algorithm then proceeds on the remaining set. Each subsequent iteration thus decreases the number of points by a constant factor. After $O(\log n)$ iterations all points of $\tilde{P}$ must be captured. □

Claim 3.4. *The data structure for the subspace $\tilde{U}$ that captures $\tilde{p}^*$ always reports this point as the $(1+\epsilon/4)$-approximate nearest neighbor of $\tilde{q}$ (in $\tilde{U}$).*

We prove Claim 3.4 in Appendix A. The proof has to overcome the disparity between projection onto U, in which p^* is the nearest neighbor, and onto the subspace $\tilde{U}$ used by the algorithm. We achieve this by careful applications of the triangle inequality and Pythagoras' Theorem, using the bounds on the noise-magnitude bound $\alpha < \epsilon/16$ and on the distance to the subspace $\sqrt{2}\alpha$.

We can now complete the proof of Theorem 3.1. By Claim 3.4, $\tilde{p}^*$ is always reported by the k-dimensional data structure it is assigned to. But this is the closest point overall, by Claim 3.2, and thus our algorithm eventually reports this point $\tilde{p}^*$, which proves the correctness part of Theorem 3.1. To argue the time and space guarantees, we just note that computing one PCA on n points takes time $O(n + d^3)$, and there are in total $O(\log n)$ PCAs to compute, and obviously also k-dimensional NNS data structures to query against.

4 Stability of a Top PCA Subspace

Before continuing to the full iterative-PCA algorithm, we need to address the challenge of controlling the stability of the PCA subspace under random noise. In particular, we will need to show that the PCA subspace $\tilde{U}$ computed from the noisy dataset $\tilde{P}$ is "close" to the original subspace U. We establish this rigorously using the sine-theta machinery developed by Davis and Kahan [DK70] and by Wedin [Wed72].

Notation Throughout, $s_j(M)$ denotes the j-th largest singular value of a real matrix M, and $\|M\| = s_1(M)$ denotes its spectral norm, while $\|M\|_F$ denotes the Frobenius norm of M. All vector norms, i.e. $\|v\|$ for $v \in \mathbb{R}^d$, refer to the ℓ_2-norm. We provide a more self-contained quick review of basic matrix analysis and spectral properties in Appendix B.

4.1 Wedin's $\sin\theta$ Theorem

The $\sin\theta$ *distance* between two subspaces B and A of $\mathbb{R}^d$ is defined as

$$\sin\theta(B, A) = \max_{x \in B, \|x\|=1} \min_{y \in A} \|x - y\|.$$

Observe that the minimum here is just the distance to a subspace $\mathrm{dist}(x, A)$, and it is attained by orthogonal projection. Thus, for all $x' \in B$ (not necessarily of unit length) $\mathrm{dist}(x', A) = \|x'\| \cdot \mathrm{dist}\left(\frac{x'}{\|x'\|}, A\right) \leq \|x'\| \cdot \sin\theta(B, A)$.

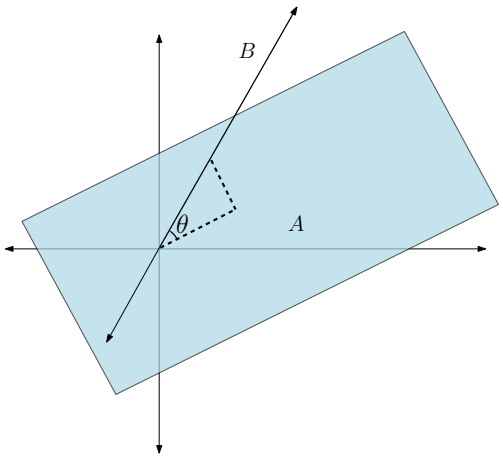

For a matrix $X \in \mathbb{R}^{n\times d}$ and an integer $m \in \{1,\ldots,d\}$, let $R_m(X)$ (resp. $L_m(X)$) denote the matrix formed by the top m right (resp. left) singular vectors of X taken in column (resp. row) order, and define $SR_m(X)$ (resp. $SL_m(X)$) as the subspace whose basis is these right (resp. left) singular vectors.

Now consider a matrix $X \in \mathbb{R}^{n\times d}$, and add to it a "perturbation" matrix $Y \in \mathbb{R}^{n\times d}$, writing $Z = X + Y$. The theorem below bounds the $\sin\theta$ distance between the top singular spaces before and after the perturbation, namely the subspaces $SR_m(Z)$ and $SR_k(X)$ for some dimensions m and k, in terms of two quantities:

1. The projection of the perturbation Y on $SR_m(Z)$ and on $SL_m(Z)$. Let $Y_R = \|YR_m(Z)\|$ and $Y_L = \|L_m(Z)Y^{\mathrm{T}}\|$.
2. The gap between the top m singular values of Z and the bottom $d-k$ singular values of X. Formally, define $\gamma = s_m(Z) - s_{k+1}(X)$.

Theorem 4.1 (Wedin's $\sin\theta$ Theorem [Wed72])**.** *In the above setting, if $m \le k \le d$ and $\gamma > 0$, then*

$$\sin\theta(SR_m(Z), SR_k(X)) \le \frac{\max\{Y_R, Y_L\}}{\gamma}.$$

4.2 Instantiating the $\sin\theta$ Theorem

We now apply the $\sin\theta$-Theorem to our semi-random model from Section 2. Let $X \in \mathbb{R}^{n\times d}$ be the matrix corresponding to our original point set P (of size $n \ge d$) lying in a subspace U of dimension $k \le d$. Let $T \in \mathbb{R}^{n\times d}$ be a perturbation matrix (noise), and then $\tilde{X} = X + T$ corresponds to our perturbed point set $\tilde{P}$. Our next theorem uses $\|T\|$ directly without assuming anything about its entries, although in our context where the entries of T are drawn from independent Gaussians of magnitude σ, Theorem B.2 implies that w.h.p. $\|T\| \le O(\sigma\sqrt{n+d})$. In fact, if the matrix T is random, then m (and possibly also γ) should be interpreted as random variables that depend on T.

Theorem 4.2. *Let $\tilde{X} = X + T$ be defined as above, and fix a threshold $\gamma_1 > 0$. If $m \le k$ is such that at least m singular values of $\tilde{X}$ are at least γ_1, then*

$$\sin\theta(SR_m(\tilde{X}), SR_k(X)) \le \frac{\|T\|}{\gamma_1},$$

where $SR_k(M)$ denotes, as before, the span of the top k right-singular vectors of a matrix M.

Proof. Towards applying Theorem 4.1, define $T_R = \|TR_m(\tilde{X})\|$ and $T_L = \|L_m(\tilde{X})T^{\mathrm{T}}\|$. The columns of $R_m(\tilde{X})$ being orthonormal implies $\|R_m(\tilde{X})\| \le 1$, and now by Fact B.1, $T_R = \|TR_m(\tilde{X})\| \le \|T\|$. We can bound also T_L similarly. Recalling Fact B.2, X has at most k non-zero singular values because the point set P lies in a k-dimensional subspace, hence the gap is $\gamma = s_m(\tilde{X}) - 0 \ge \gamma_1$. Plugging this into the $\sin\theta$ Theorem yields the bound $\sin\theta(SR_m(\tilde{X}), SR_k(X)) \le \|T\|/\gamma \le \|T\|/\gamma_1$. □

5 Iterative PCA Algorithm

We now present the iterative PCA algorithm, that solves the NNS problem for the semi-random model from Section 2. In particular, the underlying pointset lives in a k-dimensional space, but each point is also added a Gaussian noise $N_d(0, \sigma^2 I_d)$, which has norm potentially much larger than the distance between a query and it nearest neighbor. The algorithm reduces the setting to a classical k-dimensional NNS problem.

Theorem 5.1. *Suppose there is a $(1+\epsilon/8)$-approximate NNS data structure for n points in a k-dimensional space with query time F_{query}, space F_{space}, and preprocessing time F_{prep}. Assume the Gaussian-noise model* (2.1)-(2.3)*, with $\sigma(k^{1.5}\sqrt{\log n} + \sqrt[4]{k^3 d \log n}) < c\epsilon$ for sufficiently small constant $c > 0$.*

Then there is a data structure that preprocesses $\tilde{P}$, and on query $\tilde{q}$ returns $\tilde{p}^*$ with high probability. This data structure has query time $O((dk + F_{\text{query}})\sqrt{d \log n} + d^{O(1)})$, uses space $O(F_{\text{space}}\sqrt{d \log n} + d^{O(1)})$, and preprocessing time $O((nd^2 + d^3 + F_{\text{prep}})\sqrt{d \log n})$.

5.1 Algorithm Description

The iterative-PCA algorithm computes a collection $\mathcal{U}$ of $O(\sqrt{d \log n})$ subspaces, such that every point in the perturbed dataset $\tilde{P}$ is within squared distance $\Psi = d\sigma^2 + 0.001\epsilon^2$ of some subspace in the collection $\mathcal{U}$. For each such subspace $\tilde{U}_j^s \in \mathcal{U}$, we project onto $\tilde{U}_j^s$ the points captured by this subspace $\tilde{U}_j^s$, and construct on the resulting pointset a k-dimensional NNS data structure. We consider only singular vectors corresponding to sufficiently large singular values, which helps ensure robustness to noise. In particular, this threshold is $\delta(n) \triangleq c\epsilon\sqrt{\frac{n}{k}}$ for small constant $c \leq 0.001$. Also, the PCA space is computed on a sample of the current pointset only.

See Algorithm 5.1 for a detailed description of computing $\mathcal{U}$.

Algorithm 5.1 Iteratively locate subspaces.

Define $\Psi \triangleq d\sigma^2 + 0.001\epsilon^2$, $r \triangleq O(d^9 k^3 \frac{\log n}{\epsilon^2 \sigma^2})$, and $\delta(n) \triangleq c\epsilon\sqrt{\frac{n}{k}}$ for a small constant $c \leq 0.001$.
$j \leftarrow 0$, $\tilde{P}_0 \leftarrow \tilde{P}$
while $|\tilde{P}_j| > r$ **do**
  Sample r points from $\tilde{P}_j$ (with repetition) to form the set/matrix $\tilde{P}_j^s$
  $m \leftarrow$ number of singular values of $\tilde{P}_j^s$ that are at least $\delta(r)$
  $\tilde{U}_j^s \leftarrow$ the subspace spanned by the m top singular vectors of $\tilde{P}_j^s$
  $M_j \leftarrow$ all $\tilde{p} \in \tilde{P}_j \setminus \tilde{P}_j^s$ at distance $\text{dist}(\tilde{p}, \tilde{U}_j^s) \leq \sqrt{\Psi}$
  $\tilde{P}_{j+1} \leftarrow \tilde{P}_j \setminus (M_j \cup \tilde{P}_j^s)$
  $j \leftarrow j + 1$
end while
return the subspaces $\tilde{\mathcal{U}} = \{\tilde{U}_0^s, \ldots, \tilde{U}_{j-1}^s\}$, their pointsets $\{M_0, M_1, \ldots, M_{j-1}\}$, and the remaining set $R = \tilde{P}_j \bigcup \cup_{l=0}^{j-1} \tilde{P}_l^s$.

We now present the overall NNS algorithm in detail. The preprocessing stage runs Algorithm 5.1 on the pointset $\tilde{P}$, stores its output, and constructs a k-dimensional NNS data structure for each of the pointsets $M_0, \ldots, M_{j-1}$ (here j refers to the final value of this variable). Note that we also have a "left-over" set $R = \tilde{P}_j \bigcup \cup_{l=0}^{j-1} \tilde{P}_l^s$, which includes the points remaining at the end plus the sampled points used to construct the subspaces $\mathcal{U}$.

The query stage uses those j data structures to compute a $(1 + \epsilon/8)$-approximates NNS of q in each of $M_0, \ldots, M_{j-1}$, and additionally finds the NNS of q inside R by exhaustive search. It finally reports the closest point found. In the rest of this section, we will analyze this algorithm, thus proving Theorem 5.1.

We make henceforth three assumptions that hold without loss of generality. First, we assume that $\|p_i\| \geq 1$, which is without loss of generality as we can always move the pointset away from the origin. Overall, this ensures that $\|P\|_F^2 \geq |P|$.

Second, we assume that all points $\tilde{P}$ have norm at most $L \triangleq d^{3/2}$, which follows by applying a standard transformation of partitioning the dataset by a randomly shifted grid with side-length d. This transformation ensures that the query and the nearest neighbor, at distance $O(\sigma\sqrt{d})$ are in the same grid cell with probability at least $1 - o(1)$ (see, e.g., [AKS13]).

Third, we assume that $\sigma \gg \epsilon/\sqrt{d}$, as otherwise we can apply the algorithm from Section 3 directly. (The algorithm in the current section works also for small σ, but the exposition becomes simpler if we assume a larger σ.) In the context of our model of Section 2 and Lemma 2.3, this is equivalent to asserting $d \gg \log n$.

Finally, we remark that the algorithm can be changed to not use explicitly the value of σ, by taking only the closest $O\left(\sqrt{\frac{\log n}{d}}\right)$ fraction of points to a given space $\tilde{U}_j^s$. We omit the details.

5.2 Analysis

We now present a high-level overview of the proof. First, we characterize the space $\tilde{U}_j^s$, and in particular show that it is close to (a subspace of) the original space U, using the sine-theta machinery and matrix concentration bounds. Second, we use the closeness of $\tilde{U}_j^s$ to U to argue that: (a) projection of the noise onto $\tilde{U}_j^s$ is small; and (b) the projection of a point $\tilde{p}$ is approximately $\|p\|$, *on average*. Third, we use these bounds to show that the space $\tilde{U}_j^s$ captures a good fraction of points to be put into M_j, thus allowing us to bound the number of iterations. Fourth, we show that, for each point $\tilde{p} = p + t$ that has been "captured" into M_j, its projection into $\tilde{U}_j^s$ is a faithful representations of p, in the sense that, for such a point, the distance to the projection of $\tilde{q}$ onto $\tilde{U}_j^s$ is close to the original distance (before noise). This will suffice to conclude that the k-dimensional NNS for that set M_j shall return the right answer (should it happen to have the nearest neighbor p^*).

Slightly abusing notation, let P represent both the pointset and the corresponding $n \times d$ matrix, and similarly for $\tilde{P}$ or a subset thereof like $\tilde{P}_j$. Let T be the noise matrix, i.e., its rows are the vectors t_i and $\tilde{P} = P + T$.

Using bounds from random matrix theory (see Lemma B.3), w.h.p. every restriction of T to a subset of at least d rows gives a submatrix T' of spectral norm $\|T'\| \leq \eta(|T'|) = O(\sigma\sqrt{|T'| \cdot \log n})$. In addition, by Corollary 2.2 and the parameters of our model in Theorem 5.1, w.h.p.

$$\forall p_i \in P, \qquad |\|t_i\|^2 - \sigma^2 d| \leq 0.0001\epsilon^2. \tag{5.1}$$

We assume in the rest of the proof that these events occur. Since both are high probability events, we may use a union bound and assume they occur over all iterations without any effect of conditioning on the points.

The thrust of our proof below is to analyze one iteration of Algorithm 5.1. We henceforth use j to denote an arbitrary iteration (not its final value), and let $n_j = |\tilde{P}_j| > r$ denote the number of points at that iteration.

5.2.1 Analysis: Characterization of the PCA space of the sampled set

Define $\tilde{U}_j^s$ and $\tilde{U}_j$ to be the PCA space of $\tilde{P}_j^s$ and $\tilde{P}_j$ respectively, i.e., full current set and sampled set. Suppose the dimension of $\tilde{U}_j^s$ and $\tilde{U}_j$ is $m \leq k$ and $m \leq \ell \leq k$ respectively where m is set according to the thresholding step in Algorithm 5.1 and ℓ will be specified later. We show that the computed PCA space $\tilde{U}_j^s$ is close to U using the sine-theta machinery established in Section 4.2. We consider the point sets as matrices, and concatenate the following two observations for deriving our result:

- The PCA space of sampled noisy set (scaled) is close to that of the full noisy set.
- The PCA space of the noisy set is close to that of the unperturbed set.

We now analyze the effects of sampling. We sample r points from the current set $\tilde{P}_j$ of size n_j. We use the following standard matrix concentration.

Theorem 5.2 (Rudelson and Vershynin [RV07], Thm 3.1)**.** *Suppose we sample (with replacement) r row vectors from an n-size set $A \subset \mathbb{R}^d$ (represented as a $n \times d$ matrix), and call them set $Y = \{y_1, \dots y_r\}$. Then, for any $t \in (0,1)$, and $L = \max_{a \in A} \|a\|$:*

$$\Pr\left[\left\|\frac{n}{r}\sum_{y\in Y} y^{\mathrm{T}}y - A^{\mathrm{T}}A\right\| > t\|A^{\mathrm{T}}A\|\right] \leq 2e^{-\Omega\left(\left(t^2/L^2\right)\cdot r/\log r\right)}.$$

Corollary 5.3. *For $t \in (0,1)$ if we sample $r = O(\log^2 n \cdot L^2/t^2)$ vectors from $\tilde{P}_j$ to obtain $\tilde{P}_j^s$, we have that w.h.p.:*

$$\|\tfrac{n_j}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s - \tilde{P}_j^{\mathrm{T}}\tilde{P}_j\| \leq L^2 n_j \cdot t.$$

Proof. Instantiate the theorem for $A = \tilde{P}_j$ to obtain $\|\frac{n_j}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s - \tilde{P}_j^{\mathrm{T}}\tilde{P}_j\| \leq t\|\tilde{P}^{\mathrm{T}}\tilde{P}\|$. We simplify the right hand side of this equation. First, by Fact B.3 we have that $t\|\tilde{P}_j^{\mathrm{T}}\tilde{P}_j\| = t\|\tilde{P}_j\|^2$. Next by Fact B.2, $t\|\tilde{P}_j\|^2 = \sum_{\tilde{p}\in\tilde{P}_j}\|\tilde{p}\|^2 t \leq n_j \max_{\tilde{p}\in\tilde{P}_j}\|\tilde{p}\|^2 t \leq L^2 n_j t$. To prove this event succeeds with high probability over all points n, we need only substitute the value of r directly into the theorem statement. □

Corollary 5.4. *We set $t = O\left(\frac{\epsilon\sigma\sqrt{\log n}}{L^2 k^{1.5}}\right)$, for which we need to sample $r = \Omega(L^6 k^3 \log n/\epsilon^2\sigma^2) = \Omega(d^9 k^3 \log n/\epsilon^2\sigma^2)$. Then we obtain:*

$$\|\tfrac{n_j}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s - \tilde{P}_j^{\mathrm{T}}\tilde{P}_j\| \leq n_j \cdot O\left(\epsilon\sigma\tfrac{\sqrt{\log n}}{k^{1.5}}\right) \ll \left(\frac{\delta(n_j)}{k}\right)^2,$$

for σ in the range given by our model of Theorem 5.1.

We now aim to show that $\sin\theta(\tilde{U}_j^s, \tilde{U}_j)$ and $\sin\theta(\tilde{U}_j, U_j)$ are small, and the triangle inequality for $\sin\theta$ will then show that $\sin\theta(\tilde{U}_j^s, U_j)$ is also small. Recall first that the sine-theta machinery of Theorem 4.2 requires a "gap" between the bottom most singular value considered in one subspace and the top most not considered in the other. We observe the following lemma:

Lemma 5.5. *There exists ℓ, where $m \leq \ell \leq k+1$, such that $s_\ell(\tilde{P}_j) - s_{\ell+1}(\tilde{P}_j) \geq \Omega\left(\frac{\delta(n_j)}{k}\right)$. Hence $s_\ell(\tilde{P}_j) - s_{k+1}(P_j) \geq \Omega\left(\frac{\delta(n_j)}{k}\right)$.*

Proof. First recall that by the threshold step of our algorithm, $s_m\left(\frac{n_j}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s\right) \geq \delta(n_j)^2$. Now by Fact B.4 and r set as in Corollary 5.4, we have that $s_m(\tilde{P}_j^{\mathrm{T}}\tilde{P}_j) \geq \delta(n_j)^2 - \delta(n_j)^2/k^2 \geq \frac{3}{4}\delta(n_j)^2$. Hence we have $s_m(\tilde{P}_j) > 3\delta(n_j)/4$. Also since P_j is drawn from a k dimensional subspace, $s_{k+1}(P_j) = 0$ and therefore $s_{k+1}(\tilde{P}_j) \leq \|P_j - \tilde{P}_j\| = O(\sigma\sqrt{n_j \log n}) < \delta(n_j)/16$ by Lemma B.3 and the parameters of our model. Now since $s_m(\tilde{P}_j) \geq 3\delta(n_j)/4$ and $s_{k+1}(\tilde{P}_j) \leq \delta(n_j)/16$, then there must exist ℓ with $m \leq \ell \leq k+1$ that satisfies the claim. See Figure 2 for illustration of the argument. □

Using now that $a^2 - b^2 \geq (a-b)^2$ for $a \geq b \geq 0$ we obtain:

Corollary 5.6. *There exists ℓ, $m \leq \ell \leq k+1$, such that $s_\ell\left(\tilde{P}_j{}^{\mathrm{T}}\tilde{P}_j\right) - s_{\ell+1}\left(\tilde{P}_j{}^{\mathrm{T}}\tilde{P}_j\right) \geq \Omega\left(\frac{\delta(n_j)^2}{k^2}\right)$. Hence $s_m\left(\frac{n_j}{r}\left(\tilde{P}_j^s\right)^{\mathrm{T}}\tilde{P}_j^s\right) - s_{\ell+1}(\tilde{P}_j^{\mathrm{T}}\tilde{P}_j) \geq \Omega\left(\frac{\delta(n_j)^2}{k^2}\right)$.*

Now crucially using this ℓ of Lemma 5.5 to define our cut off point for the singular vectors included in $\tilde{U}_j$, and instantiating the sine-theta theorem, we obtain the following:

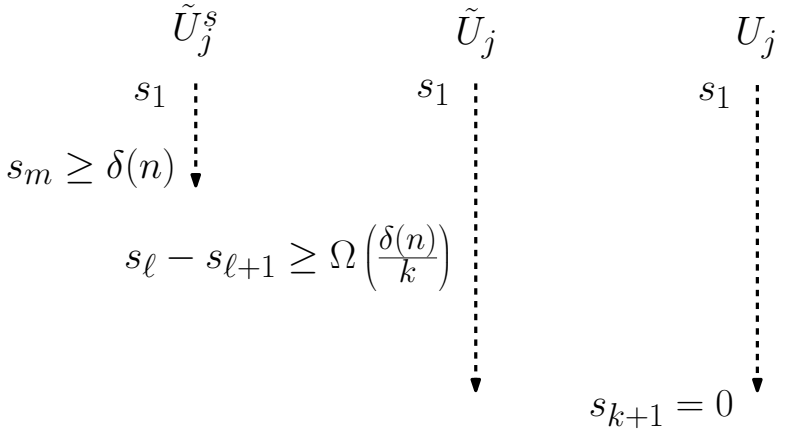

Figure 2: Take s_l as the last singular value included in $\tilde{U}_j$.

Lemma 5.7. $\sin\theta(\tilde{U}_j^s, \tilde{U}_j) \leq O\left(\frac{\sigma k^{1.5}\sqrt{\log n}}{\epsilon}\right)$ *and* $\sin\theta(\tilde{U}_j, U_j) \leq O\left(\frac{\sigma k^{1.5}\sqrt{\log n}}{\epsilon}\right)$.

Proof. First by Lemma 5.5, we have $s_l(\tilde{P}_j) - s_{k+1}(P_j) \geq O\left(\frac{\delta(n)}{k}\right)$ and recall by Lemma B.3, $\|\tilde{P}_j - P_j\| = O(\sigma\sqrt{n_j \log n})$. So instantiating the sine-theta theorem, Theorem 4.2:

$$\begin{aligned}\sin\theta(\tilde{U}_j, U_j) &= \frac{\|\tilde{P}_j - P_j\|}{s_l(\tilde{P}_j) - s_{k+1}(P_j)} \\ &= O\left(\frac{\sigma\sqrt{n_j \log n}}{\delta(n_j)/k}\right) \\ &= O\left(\frac{\sigma\sqrt{n_j \log n}}{\epsilon\sqrt{n_j}/k^{1.5}}\right) = O\left(\frac{\sigma k^{1.5}\sqrt{\log n}}{\epsilon}\right).\end{aligned}$$

Similarly for $\sin\theta(\tilde{U}_j^s, \tilde{U}_j)$, we upper bound $\|\frac{n_j}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s - \tilde{P}_j^{\mathrm{T}}\tilde{P}_j\|$ by Corollary 5.4. We lower bound the "gap" in singular values using Corollary 5.6, and hence by instantiating the sine-theta theorem:

$$\begin{aligned}\sin\theta(\tilde{U}_j^s, \tilde{U}_j) &= O\left(\frac{\epsilon n_j \sigma\sqrt{\log n}/k^{1.5}}{\delta(n_j)^2/k^2}\right) \\ &= O\left(\frac{\epsilon n_j \sigma\sqrt{\log n}/k^{1.5}}{\epsilon^2 n_j/k^3}\right) = O\left(\frac{\sigma k^{1.5}\sqrt{\log n}}{\epsilon}\right).\end{aligned}$$

□

Since $\sin\theta$ is concave in the right regime, Lemma 5.7 now gives us as a simple corollary the main claim of this subsection:

$$\sin\theta(\tilde{U}_j^s, U_j) \leq O\left(\frac{\sigma k^{1.5}\sqrt{\log n}}{\epsilon}\right). \tag{5.2}$$

5.2.2 Analysis: Noise inside the PCA space

We now show that the noise vector t_i of each point $\tilde{p}_i = p_i + t_i$ has a small component inside $\tilde{U}_j^s$. We use the $\sin\theta$ bound in Eqn. (5.2) for this. (Tighter analysis is possible directly via the randomness of the vector t_i on the first iteration, but conditioning of points selected together at each iteration of our algorithm makes this complicated at latter stages.)

Following Remark 2.1, we shall assume that noise t_i is perpendicular to U. Define $V \in \mathbb{R}^{d\times k}$ as the projection matrix onto the space U, so e.g. $t_i V$ is the zero vector, and define analogously $\tilde{V}_j^s \in \mathbb{R}^{d\times m}$ to be the projection matrix onto the m-dimensional space $\tilde{U}_j^s$.

Lemma 5.8. $\|t_i\tilde{V}_j^s\| \le \|t_i\| \cdot \sin\theta(\tilde{U}_j^s, U)$.

Proof. Let x be a unit vector in the direction of $t_i\tilde{V}_j^s$, namely, $x = \frac{t_i\tilde{V}_j^s}{\|t_i\tilde{V}_j^s\|}$. This implies $\|t_i\tilde{V}_j^s\| = x^{\mathrm{T}}t_i$. Now decompose $x \in \tilde{U}_j^s$ as $x = \sqrt{1-\beta^2}u + \beta v$ for unit vectors $u \in U$ and $v \perp U$ and some $\beta \ge 0$. Then $\beta = d(x, U) \le \sin\theta(\tilde{U}_j^s, U)$, and we conclude

$$\|t_i\tilde{V}_j^s\| = x^{\mathrm{T}}t_i = \sqrt{1-\beta^2}u^{\mathrm{T}}t_i + \beta v^{\mathrm{T}}t_i = 0 + \beta v^{\mathrm{T}}t_i \le \|t_i\| \cdot \sin\theta(\tilde{U}_j^s, U). \qquad \square$$

Corollary 5.9. *For every point p_i and iteration j, $\|t_i\tilde{V}_j^s\| \le O\left(\frac{1}{\epsilon}\sigma^2 k^{1.5}\sqrt{d\log n}\right)$.*

Proof. Substitute into Lemma 5.8 the bounds $\|t_i\| \le O(\sigma\sqrt{d})$ from Eqn. (5.1) and from Eqn. (5.2), $\sin\theta(\tilde{U}_j^s, U) \le O\left(\frac{1}{\epsilon}\sigma k^{1.5}\sqrt{\log n}\right)$. $\square$

We note that according to our model parameters in Theorem 5.1, this implies that $\|t_i\tilde{V}_j^s\| \le c\epsilon$ for a small constant c that depends only on the choice of constant in our model, and we shall use this assumption henceforth.

5.2.3 Analysis: Projection of the data into the PCA space

We now show that the component of a data point $\tilde{p}_i$ inside the PCA space $\tilde{U}_j^s$ of some iteration j, typically recovers most of the "signal", i.e., the unperturbed version p_i. More precisely, we compare the length seen inside the PCA space $\|\tilde{p}_i\tilde{V}_j^s\|$ with the original length $\|p_i\|$. While the upper bound is immediate, the lower bound holds only *on average*.

Lemma 5.10. *W.h.p., for all $\tilde{p}_i \in \tilde{P}_j$, $\|\tilde{p}_i\tilde{V}_j^s\|^2 - \|p_i\|^2 \le d\sigma^2 + 0.0001\epsilon^2$.*

Proof. Using Pythagoras' Theorem, $\|\tilde{p}_i\tilde{V}_j^s\|^2 \le \|\tilde{p}_i\|^2 = \|p_i\|^2 + \|t_i\|^2$, and the lemma follows by the noise bound (5.1). $\square$

Lemma 5.11. $\sum_{\tilde{p}_i\in\tilde{P}_j}(\|\tilde{p}_i\tilde{V}_j^s\|^2 - \|p_i\|^2) \ge -k\delta(n_j)^2$.

Proof. Let V be the projection matrix into U and P_j the non-noised version of $\tilde{P}_j$. Observe that $\tilde{P}_j V = P_j$, since the noise is orthogonal to U. Hence, by definition of PCA space (Theorem B.1):

$$\sum_{p_i\in P_j}\|p_i\|^2 = \|P_j\|_F^2 = \|\tilde{P}_j V\|_F^2 \le \sum_{l=1}^{k} s_l^2(\tilde{P}_j).$$

By Corollary 5.4, we further have that

$$\tfrac{n_j}{r}\sum_{l=1}^{k} s_l^2(\tilde{P}_j^s) \ge \sum_{l=1}^{k} s_l^2(\tilde{P}_j) - k\cdot\left(\tfrac{\delta(n_j)}{k}\right)^2 \ge \sum_i \|p_i\|^2 - k\cdot\left(\tfrac{\delta(n_j)}{k}\right)^2.$$

Or simply,

$$\frac{n_j}{r}\sum_{l=1}^{k} s_l^2(\tilde{P}_j^s) \geq \sum_i \|p_i\|^2 - k\cdot\left(\frac{\delta(n_j)}{k}\right)^2. \tag{5.3}$$

We also have:

$$\begin{aligned}\frac{n}{r}\sum_{l=1}^{m} s_l^2(\tilde{P}_j^s) &= \frac{n}{r}\|(\tilde{V}_j^s)^T(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s\tilde{V}_j^s\|_F = \|(\tilde{V}_j^s)^T\frac{n}{r}(\tilde{P}_j^s)^{\mathrm{T}}\tilde{P}_j^s\tilde{V}_j^s\|_F \\ &= \|(\tilde{V}_j^s)^T(\tilde{P}_j^{\mathrm{T}}\tilde{P}_j + Z)\tilde{V}_j^s\|_F \leq \sum_{\tilde{p}_i\in\tilde{P}_j}\|\tilde{p}_i\tilde{V}_j^s\|^2 + \|(\tilde{V}_j^s)^T Z\tilde{V}_j^s\|_F,\end{aligned}$$

where Z has spectral norm at most $\delta(n_j)^2/k^2$, and hence $\|(\tilde{V}_j^s)^T Z\tilde{V}_j^s\|_F \leq \frac{k\delta(n_j)^2}{k^2} \leq \frac{\delta(n_j)^2}{k}$.

Rearranging yields us:

$$\frac{n}{r}\sum_{l=1}^{m} s_l^2(\tilde{P}_j^s) \leq \sum_{\tilde{p}_i\in\tilde{P}_j}\|\tilde{p}_i\tilde{V}_j^s\|^2 + \frac{\delta(n_j)^2}{k}.$$

Finally, we have:

$$\begin{aligned}\frac{n_j}{r}\sum_{l=1}^{k} s_l^2(\tilde{P}_j^s) &= \frac{n_j}{r}\sum_{j=1}^{m} s_j^2(\tilde{P}_j^s) + \frac{n_j}{r}\sum_{j=m+1}^{k} s_j^2(\tilde{P}_j^s) \\ &\leq \sum_{p_i\in\tilde{P}_j}\|\tilde{p}_i\tilde{V}_j^s\|^2 + \frac{\delta(n_j)^2}{k} + \frac{n_j}{r}\sum_{j=m+1}^{k} s_j^2(\tilde{P}_j^s) \\ &\leq \sum_{p_i\in\tilde{P}_j}\|\tilde{p}_i\tilde{V}_j^s\|^2 + \frac{\delta(n_j)^2}{k} + \frac{n_j}{r}(k-1)\delta(r)^2 \\ &\leq \sum_{p_i\in\tilde{P}_j}\|\tilde{p}_i\tilde{V}_j^s\|^2 + \frac{\delta(n_j)^2}{k} + (k-1)\delta(n_j)^2\end{aligned}$$

where we employed the threshold $s_j \leq \delta(r)$ for singular values taken by Algorithm 5.1 with $j > m$, and that $\frac{n_j}{r}\delta(r)^2 = \delta(n_j)^2$ by straightforward substitution of the formula $\delta(x) = \frac{c\epsilon\sqrt{x}}{\sqrt{k}}$. The lemma follows by combining the above with Equation 5.3. □

5.2.4 Analysis: Number of iterations

We now show that each iteration captures in M_j a good fraction of the remaining points, thereby bounding the number of iterations overall. In particular, we give a lower bound on the number of indexes i such that $\tilde{p}_i$ is close to the m dimensional PCA subspace $\tilde{U}_j^s$, using results from Section 5.2.3. Note that the square of this distance for a point $\tilde{p}_i$ is precisely $\|\tilde{p}_i\|^2 - \|\tilde{p}_i\tilde{V}\|^2$. Let X and Y be quantities according to Lemmas 5.10 and 5.11, such that

$$\|\tilde{p}_i\tilde{V}_j^s\|^2 - \|p_i\|^2 \leq Y. \tag{5.4}$$

$$Xn_j \leq \sum_i(\|\tilde{p}_i\tilde{V}_j^s\|^2 - \|p_i\|^2); \tag{5.5}$$

Now let f be the fraction of i's such that $\|\tilde{p}_i\tilde{V}_j^s\|^2 - \|p_i\|^2 \leq -0.0002\epsilon^2$. Then

$$Xn_j \leq \sum_i \|\tilde{p}_i\tilde{V}_j^s\|^2 - \sum_i\|p_i\|^2 \leq (1-f)n_jY - 0.0002fn_j\epsilon^2. \tag{5.6}$$

Rearrangement of terms gives us that $f \leq \frac{Y-X}{Y+0.0002\epsilon^2}$. By Lemma 5.11 , we can set $Xn_j = -k\delta^2 = -c^2\epsilon^2 n_j = -0.00001\epsilon^2 n_j$ and so $X \leq -0.00001\epsilon^2$. And by Lemma 5.10, we have $Y \leq d\sigma^2 + 0.0001\epsilon^2$. Elementary calculations now yield

$$f \leq 1 - \Omega\left(\frac{\epsilon^2}{d\sigma^2}\right) \leq 1 - \Omega\left(\sqrt{\frac{\log n}{d}}\right).$$

(This last upper bound on f can be made tighter as dependence on σ, but we opt for the looser bound which holds in our range of parameters for simplicity.) Now for the rest of the $(1-f) \geq \Omega\left(\sqrt{\frac{\log n}{d}}\right)$ fraction of the points, the distance to the PCA subspace $\tilde{U}$ is, by Pythagoras' Theorem,

$$\|\tilde{p}\|^2 - \|\tilde{p}\tilde{V}_j^s\|^2 = \|p\|^2 + \|t\|^2 - \|\tilde{p}\tilde{V}_j^s\|^2 \leq \|t\|^2 + 0.0002\epsilon^2.$$

Since $\|t\|^2 \leq d\sigma^2 + 0.0001\epsilon^2$, we get the required inequality that a large fraction of the points is within squared distance $d\sigma^2 + 0.001\epsilon^2 = \Psi$. It follows that the fraction of points captured by $\tilde{U}_j^s$, i.e., in a single iteration, is at least $\Omega\left(\sqrt{\frac{\log n}{d}}\right)$, which immediately implies a bound on the number of iterations, as follows.

Lemma 5.12. *Algorithm 5.1 terminates in at most* $O\left(\sqrt{\frac{d}{\log n}}\log n\right) = O(\sqrt{d\log n})$ *iterations.*

5.2.5 Analysis: Correctness

It now remains to show that the data structure that captures the actual nearest neighbor $\tilde{p}^*$ will still report $\tilde{p}^*$ as the nearest neighbor to $\tilde{q}$ in the k-dimensional data structure. Suppose $\tilde{p}^*$ has been captured in j^{th} iteration, i.e., $\tilde{p}^* \in M_j$. For simplicity of exposition, let $\tilde{U} = \tilde{U}_j^s$ and $\tilde{V} = \tilde{V}_j^s$.

Note that all distance computations for a query $\tilde{q}$ are of the form $\|\tilde{q}\tilde{V} - \tilde{p}\tilde{V}\| = \|(\tilde{q}-\tilde{p})\tilde{V}\|$, where $\tilde{p}$ is a point that is close to $\tilde{U}$. Let $\tilde{q} = q + t_q$ and $\tilde{p} = p + t$. Then we have for a point $\tilde{p}$:

$$\|(\tilde{q}-\tilde{p})\tilde{V}\| = \|(q-p)\tilde{V}\| \pm O\left(\|t_q\tilde{V}\| + \|t\tilde{V}\|\right)$$

Considering the noise of the query point q, by Corollary 2.2, we have $\|t_q\tilde{V}\| \leq O(\sigma\sqrt{k} + \sigma\sqrt{\log n})$ w.h.p. Similarly, considering the noise t of a point, by Corollary 5.9, we have $\|t\tilde{V}\| \leq O(\frac{\sigma^2}{\epsilon}k^{1.5}\sqrt{d\log n})$. By the model specified in Theorem 5.1, we can set both these terms to be smaller than 0.01ϵ. Hence we have:

$$\|(\tilde{q}-\tilde{p})\tilde{V}\| = \|(q-p)\tilde{V}\| \pm 0.02\epsilon. \tag{5.7}$$

Furthermore, we have $\sin\theta(\tilde{U}, U) \leq 0.01\epsilon$. We now decompose U into two components. Let U_{in} be the projection of $\tilde{U}$ onto U, and U_{out} be the space orthogonal to U_{in} but lying in U. See Figure 5.2.5. We note that U_{out} is also orthogonal to $\tilde{U}$: otherwise some component of U_{out} would lie in the projection of $\tilde{U}$ onto U, which is a contradiction. Let V_{in} and V_{out} be the corresponding projection matrices. We likewise decompose each point p as $p_{in} \in U_{in}$ and $p_{out} \in U_{out}$.

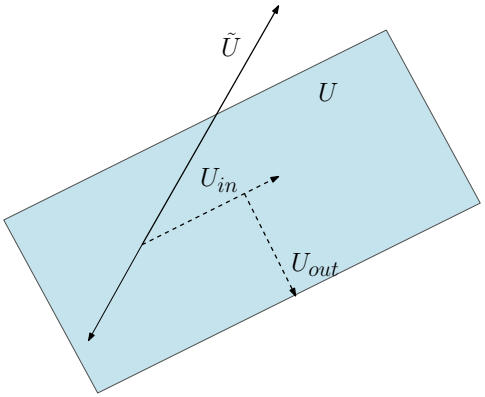

Figure 3: U_{in} is projection of $\tilde{U}$ into U, and U_{out} is the orthogonal complement of U_{in} in U.

We make the following claim, which shows that bounding $\|p_{out}\|$ for all points in M_j suffices for correctness of our algorithm.

Lemma 5.13. *If $\|p_{out}\| \leq 0.1\epsilon$ for all p captured by subspace $\tilde{U}$ (also capturing p^*), then $\tilde{p}^*$ remains a nearest neighbor to $\tilde{q}$ after projection to $\tilde{U}$. Also for any $p' \neq p^*$, we have that $\|(\tilde{q}-\tilde{p}')\tilde{V}\| \geq \left(1+\frac{\epsilon}{8}\right)\|(\tilde{q}-\tilde{p}^*)\tilde{V}\|$.*

Proof. Consider Equation 5.7.

$$\begin{aligned}
\|(\tilde{q}-\tilde{p})\tilde{V}\| &= \|(q-p)\tilde{V}\| \pm 0.02\epsilon \\
&= \|(q_{in}-p_{in}+q_{out}-p_{out})\tilde{V}\| \pm 0.02\epsilon \quad \text{(Decomposing } q, p \text{ into components in } U_{in} \text{ and } U_{out}) \\
&= \|(q_{in}-p_{in})\tilde{V}+(q_{out}-p_{out})\tilde{V}\| \pm 0.02\epsilon \\
&= \|(q_{in}-p_{in})\tilde{V}\| \pm 0.02\epsilon \quad \text{(Since } U_{out} \perp \tilde{U}) \\
&= (1 \pm 0.01\epsilon)\|q_{in}-p_{in}\| \pm 0.02\epsilon \quad \text{(Since } \sin\theta(U_{in}, \tilde{U}) \leq 0.01\epsilon)
\end{aligned}$$

To summarize these last calculations, we have:

$$\|(\tilde{q}-\tilde{p})\tilde{V}\| = (1 \pm 0.01\epsilon)\|q_{in}-p_{in}\| \pm 0.02\epsilon. \tag{5.8}$$

Next note by Pythagoras:

$$\|q_{in}-p_{in}\|^2 = \|q-p_{in}\|^2 - \|q_{out}\|^2. \tag{5.9}$$

Also observe from the triangle inequality and the assumption of our lemma that $\|p_{out}\| \leq 0.1\epsilon$, $\forall p \in P$ captured by the subspace, $\|q-p_{in}\| = \|q-p\| \pm \|p_{out}\| = \|q-p\| \pm 0.1\epsilon$.

Hence, if p^* is captured by the data structure, $\|q-p^*_{in}\|^2 = (\|q-p^*\| \pm 0.1\epsilon)^2 \leq 1+0.25\epsilon$. Similarly for $p' \neq p^*$, we have $\|q-p'_{in}\|^2 \geq (1+\epsilon-0.1\epsilon)^2 \geq 1+1.8\epsilon$. This gives:

$$\frac{1+1.8\epsilon}{1+0.25\epsilon} \leq \frac{\|q-p'_{in}\|^2}{\|q-p^*_{in}\|^2} \leq \frac{\|q-p'_{in}\|^2 - \|q_{out}\|^2}{\|q-p^*_{in}\|^2 - \|q_{out}\|^2} = \frac{\|q_{in}-p'_{in}\|^2}{\|q_{in}-p^*_{in}\|^2},$$

where we crucially used in the second step the fact that subtracting the same quantity from both numerator and denominator of a fraction can only increase it. Some elementary algebraic manipulation shows then:

$$\frac{\|q_{in}-p'_{in}\|}{\|q_{in}-p^*_{in}\|} \geq 1+\frac{\epsilon}{4}. \tag{5.10}$$

Now we lower bound $\|q_{in} - p'_{in}\|$. We do so as follows:

$$\begin{aligned}\|q_{in} - p'_{in}\|^2 - \|q_{in} - p^*_{in}\|^2 =& \|q - p'_{in}\|^2 - \|q - p^*_{in}\|^2 && \text{(By Equation 5.9)}\\ \geq& (1 + 0.9\epsilon)^2 - (1 + 0.1\epsilon)^2 \geq 1.6\epsilon.\end{aligned}$$

This implies:

$$\|q_{in} - p'_{in}\| \geq 1.2\epsilon \tag{5.11}$$

Finally, we come to our main claim of the ratio of $\|(\tilde{q} - \tilde{p}')\tilde{V}\|$ to $\|(\tilde{q} - \tilde{p}^*)\tilde{V}\|$. By Equation 5.8, this is at least:

$$\frac{\|(\tilde{q} - \tilde{p}')\tilde{V}\|}{\|(\tilde{q} - \tilde{p}^*)\tilde{V}\|} \geq \frac{(1 - 0.01\epsilon)\|q_{in} - p'_{in}\| - 0.02\epsilon}{(1 + 0.01\epsilon)\|q_{in} - p^*_{in}\| + 0.02\epsilon}.$$

Substituting the lower bound on $\frac{\|q_{in} - p'_{in}\|}{\|q_{in} - p^*_{in}\|}$ from Equation 5.10, and the lower bound on $\|q_{in} - p'_{in}\|$ from Equation 5.11 completes our claim. □

Next we derive a sufficient condition for $\|p_{out}\|$ to be bounded as desired. For a vector a, we define $a\tilde{V}$ and $a^{\perp\tilde{U}}$ to be the components lying in and orthogonal to subspace $\tilde{U}$ respectively.

The quantity of interest is c_p, defined as the cosine of the angle between $(p_{in})^{\perp\tilde{U}}$ and $t^{\perp\tilde{U}}$, for a point $\tilde{p} = p_{in} + p_{out} + t$.

Lemma 5.14. *Decompose p as $p = p_{in} + p_{out} + t$, where $t \perp U$. Suppose $c_p \leq C\frac{\epsilon}{\sigma\sqrt{d}}$ for suitable choice of constant C. Then $\|p_{out}\| \leq 0.1\epsilon$ w.h.p.*

Proof. We show first that to upper-bound p_{out}, it suffices to lower bound $\|(p_{in} + t)^{\perp\tilde{U}}\|^2$ by $d\sigma^2 - 0.009\epsilon^2$. Indeed, for captured points, we have by construction:

$$\|\tilde{p}\|^2 - \|\tilde{p}\tilde{V}\|^2 \leq \Psi \tag{5.12}$$

where $\Psi = d\sigma^2 + 0.001\epsilon^2$.

We convert the inequality to our desired condition as follows:

$$\begin{aligned}&\|\tilde{p}\|^2 - \|\tilde{p}\tilde{V}\|^2 \leq d\sigma^2 + 0.001\epsilon^2\\ &\|p_{out}\|^2 + \|p_{in} + t\|^2 - \|\tilde{p}\tilde{V}\|^2 \leq d\sigma^2 + 0.001\epsilon^2\\ &\text{(Since by Pythagoras, } \|\tilde{p}\|^2 = \|p_{out}\|^2 + \|p_{in} + t\|^2\text{)}\\ &\|p_{out}\|^2 + \|(p_{in} + t)\tilde{V}\|^2 + \|(p_{in} + t)^{\perp\tilde{U}}\|^2 - \|\tilde{p}\tilde{V}\|^2 \leq d\sigma^2 + 0.001\epsilon^2\\ &\text{(Decomposing } p_{in} + t \text{ orthogonal to and lying in subspace } \tilde{U}\text{)}\\ &\|p_{out}\|^2 + \|\tilde{p}\tilde{V}\|^2 + \|(p_{in} + t)^{\perp\tilde{U}}\|^2 - \|\tilde{p}\tilde{V}\|^2 \leq d\sigma^2 + 0.001\epsilon^2\\ &\text{(Since } p_{out} \perp \tilde{V} \text{ by construction, hence } \tilde{p}\tilde{V} = (p_{in} + t)\tilde{V}\text{)}\\ &\|p_{out}\|^2 + \|(p_{in} + t)^{\perp\tilde{U}}\|^2 \leq d\sigma^2 + 0.001\epsilon^2\\ &\|p_{out}\|^2 \leq 0.001\epsilon^2 + d\sigma^2 - \|(p_{in} + t)^{\perp\tilde{U}}\|^2\end{aligned}$$

Clearly now if $\|(p_{in} + t)^{\perp\tilde{U}}\|^2 \geq d\sigma^2 - 0.009\epsilon^2$, then $\|p_{out}\|^2 \leq 0.01\epsilon^2$ and would complete our proof.

We now show how the bound on c_p implies the required lower bound on $\|(p_{in} + t)^{\perp\tilde{U}}\|^2$. First note by the law of cosines, that

$$\|(p_{in} + t)^{\perp\tilde{U}}\|^2 = \|p_{in}^{\perp\tilde{U}}\|^2 + \|t^{\perp\tilde{U}}\|^2 - 2\|p_{in}^{\perp\tilde{U}}\|\|t^{\perp\tilde{U}}\|c_p. \tag{5.13}$$

Next note that $\|t^{\perp\tilde{U}}\|^2 = d\sigma^2 \pm 0.001\epsilon^2$ w.h.p. and a suitably small constant c in our bound on σ in the model parameters. This follows from decomposing $\|t\|^2 = \|t^{\perp\tilde{U}}\|^2 + \|t\tilde{V}\|^2$ by Pythagoras, the concentration on $\|t\|^2$ by Equation 5.1 and the upper bound on $\|t\tilde{V}\|^2$ by Corollary 5.9. We now solve to find the desired condition on c_p.

$$\begin{aligned}
&\|p_{in}^{\perp\tilde{U}}\|^2 + \|t^{\perp\tilde{U}}\|^2 - 2\|p_{in}^{\perp\tilde{U}}\|\|t^{\perp\tilde{U}}\|c_p \geq d\sigma^2 - 0.009\epsilon^2 \\
&\|p_{in}^{\perp\tilde{U}}\|^2 + (d\sigma^2 \pm 0.001\epsilon^2) - 2\|p_{in}^{\perp\tilde{U}}\|\|t^{\perp\tilde{U}}\|c_p \geq d\sigma^2 - 0.009\epsilon^2 \\
&\|p_{in}^{\perp\tilde{U}}\|^2 - 2\|p_{in}^{\perp\tilde{U}}\|\|t^{\perp\tilde{U}}\|c_p \geq -0.008\epsilon^2 \\
&2\|p_{in}^{\perp\tilde{U}}\|\|t^{\perp\tilde{U}}\|c_p \leq \|p_{in}^{\perp\tilde{U}}\|^2 + 0.008\epsilon^2 \\
&c_p \leq \frac{1}{2\|t^{\perp\tilde{U}}\|}\left(\|p_{in}^{\perp\tilde{U}}\| + 0.008\left(\frac{\epsilon^2}{\|p_{in}^{\perp\tilde{U}}\|}\right)\right).
\end{aligned}$$

Noting that the minimum of $x + \frac{\alpha}{x}$ is $\sqrt{\alpha}$ for any fixed α and that $\|t^{\perp\tilde{U}}\|$ is $O(\sqrt{d}\sigma)$ we obtain that $c_p \leq C\frac{\epsilon}{\sqrt{d}\sigma}$ is a sufficient constraint for suitable choice of constant C. □

The final component is to prove that c_p is indeed small. Since the vector t is random, this is generally not an issue: $t^{\tilde{U}}$ and $(p_{in})^{\tilde{U}}$ will be independent for all points $\tilde{p} \in \tilde{P}_j \setminus \tilde{P}_j^s$ (but not for the sampled points – this is the reason they are never included in M_j). However, in subsequent iterations, this may introduce some conditioning on t_i, so we need to be careful and argue about all iteration "at once". In fact, we show a result slightly stronger than that required by Lemma 5.14 (within the parameters of our model):

Lemma 5.15. *Fix a "horizon" j, and a point $\tilde{p} = p + t$ that has not been sampled into any set $\tilde{P}_l^s$ for all iterations $l \leq j$ ($\tilde{p}$ may or may not have been captured at some iteration $l \leq j$). Then $c_p \leq O\left(\frac{\sqrt{\log n}}{\sqrt{d}}\right)$ at iteration j, with high probability.*

Proof. The assumption on not having been sampled means that we can think of running the algorithm, disallowing $\tilde{p}$ to be in the sample. In this situation, we can run the entire algorithm, up to iteration j, without $\tilde{p}$ — it does not affect the rest of the points in any way. Hence, the algorithm can sample sets $\tilde{P}_0^s, \ldots \tilde{P}_j^s$, and fix spaces $\tilde{U}_0^s, \ldots \tilde{U}_j^s$, without access to $\tilde{p}$. Furthermore, for each $l = 0 \ldots j$, the vector $p_{in}^{\perp\tilde{U}_l^s}$ is some vector, *independent* of the noise t. Hence we can "reveal" the vector t, after having fixed all vectors $p_{in}^{\perp\tilde{U}_l^s}$, for $l = 0 \ldots j$. The vector t will have angle $O(\frac{\sqrt{\log n}}{\sqrt{d}})$ with all of them, with high probability. Note that, at this moment, it does not actually matter whether $\tilde{p}$ was captured early on or not. (Again, t is admittedly conditioned via bounds on $\|T\|$ and $\|t\|$, but since these are "whp" events, they do not affect the conclusion.) □

5.2.6 Algorithm performance

We now remark on the resulting parameters of the algorithm.

Processing an iteration of the preprocessing stage takes $O(rd^2 + d^3 + ndk) = O(nd^2)$ time for: computing $\tilde{P}_j^s$, the PCA space, and M_j respectively. Hence, over $O(\sqrt{d\log n})$ iterations, together with preprocessing of the k-dimensional NNS data structures, we get preprocessing time $O((nd^2 + d^3 + F_{\text{prep}})\sqrt{d\log n})$.

Space requirement is essentially that of $O(\sqrt{d\log n})$ instances of k-dimensional NNS data structure, plus the space to store $O(\sqrt{d\log n})$ spaces $\tilde{U}_j^s$, and the left-over set R.

The query time is composed of: computing the projections into $O(\sqrt{d\log n})$ subspaces, querying the k-dimensional NNS data structures, and computing the distances to left-over points in R. Overall, this comes out to $O(dk \cdot \sqrt{d\log n} + \sqrt{d\log n} \cdot F_{\text{query}} + d|R|)$.

6 PCA tree

We now present our second spectral algorithm, which is closely related to the PCA tree [Spr91, VKD09]. We first give the algorithm and then present its analysis. Overall, we prove the following theorem.

Theorem 6.1. *Consider the Gaussian-noise model* (2.1)-(2.3)*, and assume its parameters satisfy* $\sigma < \kappa \cdot \min\left\{\frac{\epsilon}{\sqrt{k\log n}}, \frac{\epsilon}{\sqrt{k}\sqrt[4]{d\log n}}\right\}$*, for sufficiently small constant* $\kappa > 0$*. There exists a data structure that preprocesses* $\tilde{P}$*, and then given the query* $\tilde{q}$*, returns the nearest neighbor* $\tilde{p}^*$ *w.h.p.*[5] *And w.h.p. the query time is* $(k/\epsilon)^{O(k)} \cdot d^2$*, the space requirement is* $O(nd)$*, and the preprocessing time is* $O(n^2 d + nd^3)$*.*

The algorithm itself is deterministic.

6.1 Algorithm description

The algorithm constructs one-space partitioning tree hierarchically, where each tree node is associated with a subset of the pointset $\tilde{P}$. We start with the root of the tree, associated with all n points $\tilde{P}$. Now at each tree node x, we take the pointset associated with x, termed $\tilde{P}_x^{in}$. First, we perform a process called "de-clumping", which just discards part of the dataset, to obtain a set $\tilde{P}_x \subseteq \tilde{P}_x^{in}$. We describe this process at the end.

The main operation at a node is to take the top centered-PCA direction of $\tilde{P}_x$, termed v_x. By centered-PCA we mean subtracting from each vector in $\tilde{P}_x$ their average $a = \frac{1}{|\tilde{P}_x|}\sum_{\tilde{p}\in\tilde{P}_x}\tilde{p}$, and then taking the top PCA direction. Now, let $\theta \triangleq \frac{\epsilon}{1000k^{3/2}}$ and let Θ be the partition of the real line into segments of length θ, namely $\Theta = \{[\theta i, \theta(i+1)) \mid i \in \mathbb{Z}\}$. Then we partition $\tilde{P}_x$ into parts depending on which segment from Θ the projection of a point $\tilde{p} \in \tilde{P}_x$ onto v_x falls into. Then, we orthogonalize with respect to v_x, namely, transform each point $\tilde{p} \in \tilde{P}_x$ into $\tilde{p}' = \tilde{p} - \langle \tilde{p}, v_x\rangle v_x$. For each non-empty segment of Θ we produce a child of x associated with the points that fall into that segment, and repeat recursively on it. We stop once the current tree node has at most d points associated with it.

During a query, we follow the tree into all the buckets (slabs) that intersect a ball of radius $1+\epsilon/2$ around $\tilde{q}$. In each leaf, compute the exact distance from q to all points associated to that leaf. Finally, report the closest point found.

We now describe the de-clumping procedure that is done at each node. We compute the top centered-singular value of $\tilde{P}_x^{in}$. If this value is at least $\lambda_c = \lambda_c(|\tilde{P}_x^{in}|) \triangleq \frac{\epsilon}{16}\sqrt{|\tilde{P}_x^{in}|/k}$, then set $\tilde{P}_x \triangleq \tilde{P}_x^{in}$. Otherwise, find the closest pair of points in $\tilde{P}_x^{in}$, and let δ denote their squared-distance. Remove all the pairs of points in $\tilde{P}_x^{in}$ that have squared-distance at most $\delta + \epsilon^2/2$, to obtain $\tilde{P}_x$. (The removal is iterative, proceeding in arbitrary order.)

6.2 Analysis: Tree Depth

The key to the analysis is to show that our PCA tree has depth at most $2k$. The rest of analysis will follow as we show in later sections.

In the analysis, we use the centered-PCA directions. For this purpose, we first define the centering operation $c(A)$ for a set/matrix of points: $c(A) \triangleq A - \frac{1}{|A|}\sum_{p\in A} p$. Then the centered singular value, denoted $\|A\|_c$, is $\|c(A)\|$. Note that the norm still satisfies the triangle inequality.

Lemma 6.2 (Tree Depth)**.** *The constructed PCA tree has depth at most* $2k$*.*

[5] The probability is over the randomness from the model.

Proof. We first analyze the effect of orthogonalization on the points $\tilde{p}$. Fix some node x at a level $1 \leq l \leq 2k$, and some point $\tilde{p} = p + t_p$ reaching it. Call $\tilde{p}^x \in \tilde{P}_x^{in}$ as its version at node x, after the anterior orthogonalizations at the ancestor nodes. Also, define $n_x \triangleq |\tilde{P}_x|$.

We view each step of orthogonalization as two displacement processes. If we have orthogonalized the point with respect to some vector v, this is equivalent to snapping (projecting) the point $\tilde{p}^x$ to the hyperplane defined by $\{z \in \mathbb{R}^d \mid zv = \theta \cdot \lfloor \frac{\tilde{p}^x \cdot v}{\theta} \rfloor\}$, and then moving the hyperplane towards the origin. Most importantly, all points from node x going to the same child will be all snapped to the same hyperplane. The snapping to the hyperplane moves the point $\tilde{p}^x$ by a vector of norm at most θ. Note that, while the algorithm also moves the hyperplane to pass through the origin, this does not change the relative distances of the points in this hyperplane.

Thus we can write each point $\tilde{p}^x$ reaching node x as $\tilde{p}^x = \tilde{p} + m^x + m_p^x$, where m^x is the sum of all the hyperplane moves (and is dependent on the node x only), and m_p^x which is the sum of the "snapping moves" and depends on the actual point. We observe that m_p^x has small norm, and, in particular $\|m_p^x\|^2 \leq l \cdot \theta^2 \leq 2k\theta^2$, since each move is in an orthogonal direction with respect to the previous moves.

Below we assume that Lemma B.3 holds. Also, for any two points p_1, p_2, the norm of difference of the noises is $\|t_{p_1} - t_{p_2}\|^2 = 2\sigma^2 d \pm 0.1\epsilon^2$ according to Corollary 2.2 for $\sigma \ll \epsilon / \sqrt[4]{d \log n}$.

The main part of the proof is to prove that the top PCA direction v_x at each node x is close to U. We prove this by induction over levels.

Claim 6.3 (Induction hypothesis)**.** *Consider some node x at level l, which contains at least $d = \Omega(\log n)$ points. Let v_x be the top centered-PCA direction of $\tilde{P}_x$. The projection of v_x onto $U^\perp$ is at most $\gamma = O(\sigma\sqrt{\log n} \cdot \sqrt{k}/\epsilon)$.*

Before proving the induction hypothesis, we need to show an additional claim, which characterizes the result of de-clumping in the current node x: essentially that the top PCA direction is heavy in the space U. For the claim below, we assume that Claim 6.3 is true at all levels *above* the current node x.

For a node x, we define helper sets/matrices P_x, M_x, T_x as follows. First, consider points $\tilde{P}_x$, take their non-noised versions living in U (as in the model), and move using the vector m^x; this gives the set P_x. Also, let T_x be the noise vectors of $\tilde{P}_x$. Define matrix M_x as being composed of movements m_p^x for all points p in $\tilde{P}_x$. Note that $\tilde{P}_x = P_x + M_x + T_x$, and that $\|T_x\|_c \leq \|T_x\| \leq \eta(n_x) \leq \lambda_c(n_x)$, where $\eta(n_x)$ is the function from Lemma B.3.

Claim 6.4. *Suppose we performed the de-clumping step on $\tilde{P}_x^{in}$, to obtain $\tilde{P}_x$. For v_x the top centered-PCA direction of $\tilde{P}_x$, we have that $\|c(P_x)v_x\| \geq \lambda_c - \eta(n_x)$.*

Proof. Suppose the top singular value of $c(\tilde{P}_x^{in})$ is at least λ_c (in which case no de-clumping is done and $\tilde{P}_x = \tilde{P}_x^{in}$). Hence, considering $P_x = \tilde{P}_x - (\tilde{P}_x - P_x)$, which also implies $c(P_x) = c(\tilde{P}_x) - c(\tilde{P}_x - P_x)$, we have

$$\|c(P_x)v_x\| \geq \|c(\tilde{P}_x)v_x\| - \|c(\tilde{P}_x - P_x)v_x\| \geq \lambda_c - \eta(|\tilde{P}_x|),$$

since $\|c(\tilde{P}_x - P_x)v_x\| = \|c(T_x)v_x\| \leq \eta(|\tilde{P}_x|)$ by Lemma B.3, and the fact that $M_x v_x = 0$.

Otherwise, the algorithm throws out some points from $\tilde{P}_x^{in}$. Define P_x^{in} similarly to P_x: take original (non-moved, no noise) versions of the points from $\tilde{P}_x^{in}$, plus the overall movement m^x. In this situation, there must be two points $p_1, p_2 \in P_x^{in}$ such that $\|p_1 - p_2\| \leq \epsilon/4$: otherwise, the top singular value of $\tilde{P}_x^{in} = P_x^{in} - (P_x^{in} - \tilde{P}_x^{in})$ would be at least $\|P_x^{in}\|_c - \eta(|\tilde{P}_x^{in}|) \geq \frac{\|p_1 - p_2\|}{2} \cdot \sqrt{|\tilde{P}_x^{in}|/k} - \eta(|\tilde{P}_x^{in}|) \geq \lambda_c$, a contradiction.

We now want to characterize the minimal distance δ in $\tilde{P}_x^{in}$. Note that the projection of $m_{p_1}^x$ and $m_{p_2}^x$ into $U^\perp$ is at most $2k\theta\gamma$, since each of the basis vectors of m_p^x has projection into $U^\perp$ at most γ. Hence, the

square of the component of $\tilde{p}_1 - \tilde{p}_2$ in $U^\perp$ is equal to:

$$(\|t_{p_1} - t_{p_2}\| \pm 2k\theta\gamma)^2 = 2\sigma^2 d \pm 0.1\epsilon^2 \pm (2k\theta\gamma)^2 \pm 5\sigma\sqrt{d} \cdot 2k\theta\gamma = 2\sigma^2 d \pm 0.1\epsilon^2 \pm 0.01\epsilon^2 \pm 0.01\epsilon^2,$$

for $\sigma \ll \epsilon/\sqrt{\log n}$ and $\sigma \ll \frac{\epsilon}{\sqrt[4]{d \log n}}$. Thus, for $\tilde{p}_1 = p_1 + t_1$ and $\tilde{p}_2 = p_2 + t_2$:

$$\delta = \|\tilde{p}_1 - \tilde{p}_2\|^2 \geq (\|t_1 - t_2\| - 2k\theta\gamma)^2 \geq 2\sigma^2 d - 0.12\epsilon^2. \tag{6.1}$$

After the de-clumping, the distance between any two distinct points $p', p'' \in P_x$, with noise vectors t', t'' respectively, must satisfy:

$$(\|p' - p''\| + 2\sqrt{2k}\theta)^2 \geq \delta + \epsilon^2/2 - (\|t' - t''\| + 2k\theta\gamma)^2 \geq \epsilon^2/2 - 2 \cdot 0.12\epsilon^2 \geq \epsilon^2/4.$$

Hence $\|p' - p''\| \geq \epsilon/2 - 0.01\epsilon > \epsilon/4$, which means that $\|\tilde{P}_x\|_c \geq \lambda_c$ (as already argued above). Hence we can apply the same argument as above, this time for $\tilde{P}_x$ instead of $\tilde{P}_x^{in}$. □

We now prove the induction hypothesis, namely Claim 6.3, for the current node x.

Proof of Claim 6.3. Let $\tilde{P}_x$ be the points contained in x. By Claim 6.4, we have $\lambda_{PM} \triangleq \|P_x + M_x\|_c \geq \|c(P_x + M_x)v_x\| = \|c(P_x)v_x\| \geq \lambda_c - \eta(|\tilde{P}_x|) \geq \lambda_c/2$.

Decompose the top centered-PCA direction v_x of $\tilde{P}_x$ as $v_x = \sqrt{1-\alpha^2}u + \alpha u'$, where $u \in U$ and $u' \perp U$ are of norm 1. Note that α is exactly the projection of v_x onto $U^\perp$.

We will bound α by lower and upper bounding $\|\tilde{P}_x\|_c$ as a function of α and λ_{PM}. We start with the upper bound on $\|\tilde{P}_x\|_c$. For this we decompose the matrix $c(P_x+M_x)$ into the component in U, called $c(P_x+M_x^U)$, and the perpendicular one, called $c(M_x^\perp)$. Note that $c(P_x)$ lies inside U, despite the movement m_x, because of the centering $c(\cdot)$. We now bound the spectral norm of $\|c(M_x^\perp)\|$, using the inductive hypothesis that the projection of each of at most $2k$ basis vectors of m_p^x onto $U^\perp$ is at most γ:

$$\|c(M_x^\perp)\| \leq \max_p \sqrt{n_x}\|m_p^x\| \cdot \sqrt{2k}\gamma \leq \gamma \cdot \sqrt{n_x} 2k \cdot \theta \leq \gamma \cdot \tfrac{\epsilon}{500}\sqrt{n_x/k} \leq \tfrac{\gamma}{9} \cdot \tfrac{\epsilon}{32}\sqrt{n_x/k} \leq \tfrac{\gamma}{9}\lambda_{PM}.$$

Thus, we have that $\lambda_{PM^U} \triangleq \|P_x + M_x^U\|_c = \lambda_{PM} \pm \frac{\gamma}{9}\lambda_{PM}$.

We can now upper bound the norm of $\tilde{P}_x$:

$$\|\tilde{P}_x\|_c \leq \|c(P_x + M_x)v_x\| + \|c(T_x)v_x\| \leq \|c(P_x + M_x)v_x\| + \alpha\eta(n_x). \tag{6.2}$$

We need to bound the first term now. For this, we compute the following ratio, using that the projection of v_x into $U^\perp$ is of magnitude α:

$$\begin{aligned}
\frac{\|c(P_x + M_x)v_x\|^2}{\lambda_{PM^U}^2} &\leq \frac{\|c(P_x+M_x^U)v_x\|^2 + \|c(M_x^\perp)v_x\|^2 + 2\|c(P_x+M_x^U)v_x\|\|c(M_x^\perp)v_x\|}{\lambda_{PM^U}^2} \\
&\leq \frac{(1-\alpha^2)\lambda_{PM^U}^2 + \alpha^2 \cdot (\gamma/9)^2 \cdot \lambda_{PM}^2 + 2\alpha\sqrt{1-\alpha^2} \cdot \lambda_{PM^U} \cdot \frac{\gamma}{9}\lambda_{PM}}{\lambda_{PM^U}^2}. \\
&\leq (1-\alpha^2) + \alpha^2(\gamma/9)^2(1+\gamma/9)^2 + 2\alpha\tfrac{\gamma}{9} \cdot (1+\gamma/9) \\
&\leq 1 - \alpha^2/2 + \alpha\gamma/3.
\end{aligned} \tag{6.3}$$

On the other hand, we want to prove a lower bound on $\|\tilde{P}_x\|_c$. We define u_{PM^U} to be the centered-PCA direction of $P_x + M_x^U$: $\lambda_{PM^U} = \|(P_x + M_x^U)u_{PM^U}\|_c$. Remember that u_{PM^U} lies inside U.

$$\|\tilde{P}_x\|_c \geq \|c(\tilde{P}_x)u_{PM^U}\| = \|c(P_x + M_x^U)u_{PM^U}\| = \lambda_{PM^U}. \tag{6.4}$$

Putting together the upper bound (6.2), (6.3) and the lower bound (6.4), we obtain

$$\lambda_{PM^U} \leq \lambda_{PM^U} \sqrt{1 - \alpha^2/2 + \alpha\gamma/3} + \alpha\eta(n_x)$$

and, using that $\sqrt{1-x} \leq 1 - x/2$ for all $0 \leq x \leq 1$, we conclude:

$$\alpha \leq 4 \cdot \left(\frac{\eta(n_x)}{\lambda_{PM^U}} + \gamma/6 \right) \leq \gamma,$$

as long as $4\frac{\eta(n_x)}{\lambda_{PM^U}} < \gamma/3$. Since $\lambda_{PM^U} \geq \lambda_{PM}/2 \geq \lambda_c/4$, this is satisfied if

$$\gamma = \Theta(\sigma\sqrt{\log n}\sqrt{k}/\epsilon).$$

We are done proving the inductive hypothesis. □

The final step is to show that, because all vectors v_x along a root-to-leaf path are perpedicular to each other and are heavy in U, there cannot too many of them, and hence the tree depth is bounded.

Claim 6.5. *Fix some dimension $k > 1$ subspace $U \subset \mathbb{R}^d$. Suppose there exists $2k+1$ vectors $v_1, \ldots v_{2k+1}$, such that the projection of each into U has norm at least $1 - 1/4k$. Then at least two of the vectors are not orthogonal to each other.*

Proof. For the sake of contradiction assume that all vectors $v_1 \ldots v_{2k+1}$ are mutually orthogonal. Then let u_i be the projection of v_i into U, and let $\delta_i = v_i - u_i$. We want to bound the dot product $u_i u_j$. Consider

$$0 = v_i^{\mathrm{T}} \cdot v_j = (u_i^{\mathrm{T}} + \delta_i^{\mathrm{T}})(u_j + \delta_j) = u_i^{\mathrm{T}} u_j + \delta_i^{\mathrm{T}} \delta_j.$$

Since $|\delta_i^{\mathrm{T}} \delta_i| \leq 0.25/k$, we have that $|u_i^{\mathrm{T}} u_j| \leq 0.25/k$. Even after normalization, we have that $\left| \frac{u_i^{\mathrm{T}} u_j}{\|u_i\| \cdot \|u_j\|} \right| \leq 1/k$. Following Alon's result [Alo03], we conclude that u_i's have rank at least $(2k+1)/2 > k$, which is impossible if all u_i live in the k-dimensional space U. Thus we have reached a contradiction. □

We now combine the two claims together. Note that $\alpha \leq 1/4k$. We conclude that the height of the tree is at most $2k$, thus concluding the proof of Lemma 6.2. □

6.3 Analysis: Correctness

Now that we have established an upper bound on the tree depth, we will show that the query algorithm will indeed return the nearest neighbor $\tilde{p}^*$ for the query $\tilde{q}$ (modelled as in Section 2). We show this in two steps: first we prove the result, assuming the point $\tilde{p}^*$ was not thrown out during de-clumping. Then we show that the de-clumping indeed does not throw out the point $\tilde{p}^*$.

Lemma 6.6. *The query algorithm returns the point $\tilde{p}^*$, assuming it was not thrown out during the de-clumping process.*

Proof. Consider some non-leaf tree node x containing $\tilde{p}^*$. We need to argue that, at node x, the query algorithm follows the child where $\tilde{p}^*$ goes.

As before, let $\tilde{p}^x$ be the orthogonalized version of $\tilde{p}^*$ (above x) and m_p^x is the total amount of "hyperplane snapping" that happened to $\tilde{p}^*$. We also have that v_x has projection at most γ onto $U^\perp$ (from Claim 6.3). Hence, we have:

$$|v_x \tilde{p}^x - v_x \tilde{q}| \leq |v_x(t_{p^*} - t_q)| + |v_x(p^* - q)|,$$

again using that $v_x m_p^x = 0$. Note that, with high probability,

$$|v_x(t_{p^*} - t_q)| \leq 3\sigma\sqrt{d} \cdot \gamma \leq \epsilon/8\sqrt{k}.$$

Since this is true also for all ancestors y of x, and since all v_y, together with v_x, are mutually orthogonal, we have that:

$$\begin{aligned}
&\sum_{y \text{ ancestor of } x} |v_y(\tilde{p}^x - \tilde{q})|^2 \\
&\quad = \sum_{y \text{ ancestor of } x} |v_y(t_{p^*} - t_q) + v_y(p^* - q)|^2 \\
&\quad \leq \left(\sqrt{\sum_{y \text{ ancestor of } x} |v_y(m_p^y + |t_{p^*} - t_q|)|^2} + \sqrt{\sum_{y \text{ ancestor of } x} |v_y(p^* - q)|^2} \right)^2 \\
&\quad \leq \left(\sqrt{2k \cdot \left(\epsilon/8\sqrt{k}\right)^2} + \sqrt{\sum_{y \text{ ancestor of } x} |v_y(p^* - q)|^2} \right)^2 \\
&\quad \leq (\epsilon/4\sqrt{2} + 1)^2 \\
&\quad < \epsilon/2 + 1.
\end{aligned}$$

This means that the bucket (child node of x) of $\tilde{p}^x$ intersects a ball of radius $1 + \epsilon/2$ around $\tilde{q}$, and hence the query algorithm will go into that bucket (child). □

To complete the correctness argument, we also need to prove that p^* is never thrown out due to the de-clumping process.

Claim 6.7. *p^* is never thrown out due to the de-clumping process.*

Proof. For the sake of contradiction, suppose $\tilde{p}_x^*$ is thrown out at some node x. This means there is some point $\tilde{p}_x$ such that $\|\tilde{p}_x - \tilde{p}_x^*\|^2 \leq \delta + \epsilon^2/2$. Since $\|q - p\| - \|q - p^*\| \geq \epsilon$, we have that $\|p - p^*\| \geq \epsilon$.

We have:

$$\begin{aligned}
\delta + \epsilon^2/2 \geq \|\tilde{p}_x - \tilde{p}_x^*\|^2 &= \|p - p^* + (m_p^x - m_{p^*}^x) + (t_p - t_{p^*})\|^2 \\
&= \|(p - p^* + (m_p^x - m_{p^*}^x))U\|^2 + \|(m_p^x - m_{p^*}^x + t_p - t_{p^*})U^\perp\|^2
\end{aligned}$$

We want to bound $\|(m_p^x - m_{p^*}^x + t_p - t_{p^*})U_\perp\|^2$. Indeed, note that the projection of m_p^x onto $U^\perp$ can be at most $2k\theta \cdot \gamma \leq O(\sigma\sqrt{\log n})$, by Claim 6.3. Hence:

$$\begin{aligned}
\|(m_p^x - m_{p^*}^x + t_p - t_{p^*})U_\perp\| &\geq -O(\sigma\sqrt{\log n}) + \sqrt{2\sigma^2 d - 0.1\epsilon^2} \\
&\geq -O(\sigma\sqrt{\log n}) + \sigma\sqrt{2d} - \tfrac{0.06\epsilon^2}{\sigma\sqrt{2d}} \\
&\geq \sigma\sqrt{2d} - \tfrac{0.09\epsilon^2}{\sigma\sqrt{2d}},
\end{aligned}$$

as long as $O(\sigma\sqrt{\log n}) < \frac{0.03\epsilon^2}{\sigma\sqrt{2d}}$, which is equavalent to saying that $\sigma \leq O\left(\frac{\epsilon}{\sqrt[4]{d}\sqrt[4]{\log n}}\right)$.

Putting it all together, we have:

$$\begin{aligned}
\delta + \epsilon^2/2 \geq \|\tilde{p}_x - \tilde{p}_x^*\|^2 &\geq (\epsilon - 2 \cdot \sqrt{2k}\theta)^2 + (\sigma\sqrt{2d} - \tfrac{0.09\epsilon^2}{\sigma\sqrt{2d}})^2 \\
&\geq 0.8\epsilon^2 + 2\sigma^2 d - 0.2\epsilon^2 \\
&> 2\sigma^2 d + 0.6\epsilon^2.
\end{aligned}$$

But, as proven in Eqn. (6.1), we have that $\delta \leq 2\sigma^2 d + 0.12\epsilon^2$. We've reached a contradiction. □

6.4 Analysis: Performance

The space and preprocessing bounds follow immediately from the construction. We just need to argue about the query time.

Claim 6.8. *The query time is* $(k/\epsilon)^{O(k)}d^2$.

Proof. At each node of the tree, there are at most $O(1/\theta) = O(k^{3/2}/\epsilon)$ child nodes that are followed. Hence, in total, we reach $O(1/\theta)^{2k} = (k/\epsilon)^{O(k)}$ leaves. The factor of d^2 comes from the fact that each leaf has at most d points to check the distance against. □

References

[AI06] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 459–468, 2006.

[AINR14] A. Andoni, P. Indyk, H. Nguyen, and I. Razenshteyn. Beyond locality-sensitive hashing. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. arXiv:1306.1547.

[AIP06] A. Andoni, P. Indyk, and M. Pătraşcu. On the optimality of the dimensionality reduction method. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 449–458, 2006.

[AK01] S. Arora and R. Kannan. Learning a mixture of gaussians. *Proceedings of the Symposium on Theory of Computing (STOC)*, 2001.

[AKS13] D. Aiger, H. Kaplan, and M. Sharir. Reporting neighbors in high-dimensional euclidean spaces. In *SODA*, pages 784–803, 2013.

[Alo03] N. Alon. Problems and results in extremal combinatorics I. *Discrete Mathematics*, 273:31–53, 2003.

[AM02] S. Arya and T. Malamatos. Linear-size approximate voronoi diagrams. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 147–155, 2002.

[AMM09] S. Arya, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate nearest neighbor searching. *Journal of the ACM (JACM)*, 57(1):1, 2009.

[AMN+98] S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching. *J. ACM*, 6(45):891–923, 1998. Previously appeared in SODA'94.

[BKL06] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *23rd international conference on Machine learning*, pages 97–104. ACM, 2006. doi:10.1145/1143844.1143857.

[CG06] R. Cole and L.-A. Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *38th annual ACM symposium on Theory of computing*, pages 574–583. ACM, 2006. doi:10.1145/1132516.1132599.

[Cla88] K. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17:830–847, 1988.

[Cla94] K. Clarkson. An algorithm for approximate closest-point queries. *Proceedings of the Tenth Annual ACM Symposium on Computational Geometry*, pages 160–164, 1994.

[Das99] S. Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science*, pages 634–. IEEE Computer Society, 1999. doi:10.1109/SFFCS.1999.814639.

[DF08] S. Dasgupta and Y. Freund. Random projection trees and low dimensional manifolds. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 537–546. ACM, 2008.

[DIIM04] M. Datar, N. Immorlica, P. Indyk, and V. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. *Proceedings of the ACM Symposium on Computational Geometry (SoCG)*, 2004.

[DK70] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7:1–46, 1970. doi:10.1137/0707001.

[GK93] P. Gritzmann and V. Klee. Computational complexity of inner and outer j-radii of polytopes in finite-dimensional normed spaces. *Mathematical programming*, 59(1-3):163–213, 1993. doi:10.1007/BF01581243.

[HK13] S. Har-Peled and N. Kumar. Approximate nearest neighbor search for low-dimensional queries. *SIAM J. Comput.*, 42(1):138–159, 2013. Previously in SODA'11.

[HM05] S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *21st Annual Symposium on Computational Geometry*, pages 150–158. ACM, 2005. doi:10.1145/1064092.1064117.

[HP01] S. Har-Peled. A replacement for voronoi diagrams of near linear size. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 94–103, 2001.

[HV04] S. Har-Peled and K. R. Varadarajan. High-dimensional shape fitting in linear time. *Discrete & Computational Geometry*, 32(2):269–288, 2004. doi:10.1007/s00454-004-1118-2.

[IM98] P. Indyk and R. Motwani. Approximate nearest neighbor: towards removing the curse of dimensionality. *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 604–613, 1998.

[IN07] P. Indyk and A. Naor. Nearest neighbor preserving embeddings. *ACM Transactions on Algorithms*, 2007.

[IW03] P. Indyk and D. Woodruff. Tight lower bounds for the distinct elements problem. *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 283–290, 2003.

[JL84] W. Johnson and J. Lindenstrauss. Extensions of lipshitz mapping into hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.

[JS82] W. Johnson and G. Schechtman. Embedding ℓ_p^m into ℓ_1^n. *Acta Mathematica*, 149:71–85, 1982.

[JW13] T. S. Jayram and D. P. Woodruff. Optimal bounds for johnson-lindenstrauss transforms and streaming problems with subconstant error. *ACM Transactions on Algorithms*, 9(3):26, 2013. Previously in SODA'11.

[KL04] R. Krauthgamer and J. Lee. Navigating nets: Simple algorithms for proximity search. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2004.

[KR02] D. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. *Proceedings of the Symposium on Theory of Computing (STOC)*, 2002.

[KV09] R. Kannan and S. Vempala. Spectral algorithms. *Found. Trends Theor. Comput. Sci.*, 4(3–4):157–288, March 2009. Available from: http://dx.doi.org/10.1561/0400000025, doi:10.1561/0400000025.

[Lat05] R. Latała. Some estimates of norms of random matrices. *Proceedings of the American Mathematical Society*, 133(5):1273–1282, 2005. doi:10.1090/S0002-9939-04-07800-1.

[LM98] B. Laurent and P. Massarat. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28:1303–1338, 1998.

[McN01] J. McNames. A fast nearest-neighbor algorithm based on a principal axis search tree. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(9):964–976, 2001.

[Mei93] S. Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106:286–303, 1993.

[ML09] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *VISAPP (1)*, pages 331–340, 2009.

[MV10] A. Moitra and G. Valiant. Settling the polynomial learnability of mixtures of gaussians. In *51st Annual IEEE Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.

[NRC13] *Frontiers in Massive Data Analysis*. The National Academies Press, 2013. Available from: http://www.nap.edu/openbook.php?record_id=18374.

[Pan06] R. Panigrahy. Entropy-based nearest neighbor algorithm in high dimensions. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.

[RV07] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM (JACM)*, 54(4):21, 2007.

[RV10] M. Rudelson and R. Vershynin. Non-asymptotic theory of random matrices: extreme singular values. In *Proceedings of the International Congress of Mathematicians. Volume III*, pages 1576–1602, New Delhi, 2010. Hindustan Book Agency. arXiv:1003.2990.

[SAH08] C. Silpa-Anan and R. Hartley. Optimised kd-trees for fast image descriptor matching. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[SH09] R. Salakhutdinov and G. Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009.

[Spr91] R. Sproull. Refinements to nearest-neighbor searching in k-dimensional trees. *Algorithmica*, 6:579–589, 1991. doi:10.1007/BF01759061.

[ST09] D. A. Spielman and S.-H. Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Commun. ACM*, 52(10):76–84, October 2009. doi:10.1145/1562764.1562785.

[Ste93] G. W. Stewart. On the early history of the singular value decomposition. *SIAM review*, 35(4):551–566, 1993.

[VKD09] N. Verma, S. Kpotufe, and S. Dasgupta. Which spatial partition trees are adaptive to intrinsic dimension? In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 565–574. AUAI Press, 2009.

[VVZ02] K. R. Varadarajan, S. Venkatesh, and J. Zhang. On approximating the radii of point sets in high dimensions. In *43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 561–569. IEEE, 2002. doi:10.1109/SFCS.2002.1181980.

[VW04] S. Vempala and G. Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004. Previously in FOCS'02.

[Wed72] P.-Å. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972. doi:10.1007/BF01932678.

[Woo04] D. Woodruff. Optimal space lower bounds for all frequency moments. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2004.

[WTF08] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *Advances in neural information processing systems*, pages 1753–1760, 2008. Available from: http://books.nips.cc/papers/files/nips21/NIPS2008_0806.pdf.

[YSRL11] J. Yagnik, D. Strelow, D. A. Ross, and R.-S. Lin. The power of comparative reasoning. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2431–2438. IEEE, 2011.

A Additional Material for Section 3

Proof of Claim 3.4. Fix $p \neq p^*$ that is captured by the same $\tilde{U}$, and use the triangle inequality to write $\|p - \tilde{p}_{\tilde{U}}\| \leq \|p - \tilde{p}\| + \|\tilde{p} - \tilde{p}_{\tilde{U}}\| \leq \alpha + \sqrt{2}\alpha \leq 3\alpha$. Doing similarly for p^* we get $\|p^* - \tilde{p}^*_{\tilde{U}}\| \leq 3\alpha$, and by our assumption $\|q - \tilde{q}\| \leq \alpha$. Using all of these and the triangle inequality once again, we bound (in some sense "comparing" p vs. p^*)

$$\begin{aligned}\frac{\|\tilde{q} - \tilde{p}_{\tilde{U}}\|}{\|\tilde{q} - \tilde{p}^*_{\tilde{U}}\|} &= \frac{\|q - p\| \pm 4\alpha}{\|q - p^*\| \pm 4\alpha} = \frac{\|q - p\| \pm \frac{1}{4}\epsilon}{\|q - p^*\| \pm \frac{1}{4}\epsilon} \\ &\geq \frac{\|q - p\| - \frac{1}{4}\epsilon}{\|q - p^*\| + \frac{1}{4}\epsilon} \geq \frac{\|q - p^*\| + \frac{3}{4}\epsilon}{\|q - p^*\| + \frac{1}{4}\epsilon} > 1 + \tfrac{1}{4}\epsilon.\end{aligned}$$

Aiming to "replace" $\tilde{q}$ above with $\tilde{q}_{\tilde{U}}$, its projection of $\tilde{q}$ onto $\tilde{U}$, we use Pythagoras' Theorem (recall both $\tilde{p}_{\tilde{U}}, \tilde{p}^*_{\tilde{U}} \in \tilde{U}$),

$$\frac{\|\tilde{q}_{\tilde{U}} - \tilde{p}_{\tilde{U}}\|^2}{\|\tilde{q}_{\tilde{U}} - \tilde{p}^*_{\tilde{U}}\|^2} = \frac{\|\tilde{q} - \tilde{p}_{\tilde{U}}\|^2 - \|\tilde{q} - \tilde{q}_{\tilde{U}}\|^2}{\|\tilde{q} - \tilde{p}^*_{\tilde{U}}\|^2 - \|\tilde{q} - \tilde{q}_{\tilde{U}}\|^2} > (1 + \tfrac{1}{4}\epsilon)^2,$$

where the last inequality crucially uses the fact that the numerator and the denominator contain the exact same term $\|\tilde{q} - \tilde{q}_{\tilde{U}}\|^2$. We conclude that $\tilde{p}^*$ is indeed reported by the k-dimensional data structure it is assigned to. □

B Short Review of Spectral Properties of Matrices

We review some basic definitions of spectral properties of matrices.

B.1 Spectral Norm and Principal Component Analysis

The *spectral norm* of a matrix $X \in \mathbb{R}^{n \times d}$ is defined as $\|X\| = \sup_{y \in \mathbb{R}^d : \|y\| = 1} \|Xy\|$, where all vector norms $\|\cdot\|$ refer throughout to the ℓ_2-norm. The Frobenius norm of X is defined as $\|X\|_F = (\sum_{ij} X_{ij}^2)^{1/2}$, and let X^{T} denote the transpose of X. A *singular vector* of X is a unit vector $v \in \mathbb{R}^d$ associated with a *singular value* $s \in \mathbb{R}$ and a unit vector $u \in \mathbb{R}^n$ such that $Xv = su$ and $u^{\mathrm{T}}X = sv^{\mathrm{T}}$. (We may also refer to v and u as a pair of right-singular and left-singular vectors associated with s.)

Fact B.1. *For every two real matrices X and Y of compatible dimensions (i) $\|X + Y\| \leq \|X\| + \|Y\|$; (ii) $\|XY\| \leq \|X\| \cdot \|Y\|$; and (iii) $\|X\| = \|X^{\mathrm{T}}\|$.*

We can think of i-th row in X as a point $x_i \in \mathbb{R}^d$, and define the corresponding point set $P(X) = \{x_1, \ldots, x_n\}$. Then a unit vector $y \in \mathbb{R}^d$ maximizing $\|Xy\|$ corresponds to a best-fit line for the point set $P(X)$. The PCA (and *Singular Value Decomposition (SVD)* more generally) is a classical generalization of this notion to a best-fit k-dimensional subspace for $P(X)$, as described next (see e.g. [KV09, Chapter 1]).

Theorem B.1 ([KV09])**.** *Let $X \in \mathbb{R}^{n \times d}$, and define the vectors $v_1, \ldots, v_d \in \mathbb{R}^d$ inductively by*

$$v_j = \operatorname*{arg\,max}_{\|v\| = 1;\ \forall i < j, v^{\mathrm{T}} v_i = 0} \|Xv\|$$

(where ties for any $\arg\max$ *are broken arbitrarily). Then each $V_j = \mathrm{span}\{v_1, \ldots, v_k\}$ attains the minimum of $\sum_{i=1}^n d(x_i, W)^2$ over all j-dimensional subspaces W. Furthermore, $v_1, \ldots, v_d$ are all singular vectors with corresponding singular values $s_1 = \|Xv_1\|, \ldots, s_d = \|Xv_d\|$, and clearly $s_1 \geq \ldots \geq s_d$.*

We will later need the following basic facts (see, e.g., [Ste93]). We denote the singular values of a matrix $X \in \mathbb{R}^{n\times d}$ by $s_1(X) \geq s_2(X) \geq \ldots \geq s_d(X)$.

Fact B.2. *Let $P(X) \subset \mathbb{R}^d$ be the point set corresponding to the rows of a matrix $X \in \mathbb{R}^{n\times d}$. Then*

(a). $\|X\|_F^2 = \sum_{p\in P(X)} \|p\|^2 = \sum_{i=1}^{d} s_i(X)^2$ *and* $\|X\| = s_1(X)$.

(b). $P(X)$ *lies in a subspace of dimension k if and only if* $s_{k+1}(X) = 0$.

Fact B.3. *For any matrix X, let $X^{\mathrm{T}}X$ be the* covariance *matrix of X. Then the right singular vectors of X are also right singular vectors of $X^{\mathrm{T}}X$. Also, $s_i(X^{\mathrm{T}}X) = s_i^2(X)$.*

Fact B.4. *For two matrices X and E of compatible dimensions, $|s_j(X+E) - s_j(X)| \leq \|E\|$.*

B.2 Spectral Norms of Random Matrices

In our analysis, we will also need bounds on norms of random matrices, which we derive using standard random matrix theory. We state below a bound on the spectral norm of $T \in \mathbb{R}^{n\times d}$, a matrix of iid Gaussian vectors. We also consider its restriction T_A to any subset of rows of T, and bound the spectral norm in terms of the size of the subset, $s = |A|$, expressed as a function $\eta(s) = O(\sigma\sqrt{s \cdot \log n})$.

Theorem B.2 ([Lat05, RV10])**.** *Let matrix $T \in \mathbb{R}^{n\times d}$ have entries drawn independently from $N(0,\sigma)$. Then with probability approaching 1 asymptotically as n and d increase, $\|T\| \leq 3\sigma \max\{\sqrt{n}, \sqrt{d}\}$.*

Lemma B.3. *With high probability, for every subset A of the rows of T, with $|A| \geq d$, the corresponding submatrix T_A of T, has spectral norm $\|T_A\| \leq \eta(|A|) = O(\sigma\sqrt{|A| \cdot \log n})$.*

Proof. Fix some A of size s. It is known from random matrix theory that $\|T_A\| \leq \sigma\sqrt{Cs \log n}$ with probability at least $1 - e^{-\Omega(Cs\log n)} = 1 - n^{-\Omega(Cs)}$ [RV10, Proposition 2.4]. For a large enough constant $C > 1$, since there are at most $\binom{n}{s} \leq (n/s)^{O(s)}$ such subsets A, by a union bound, with high probability none of the sets will fail. Another union bound over all sizes s completes the claim. □